arXiv:1409.6981v1 [stat.ME] 24 Sep 2014

Unsupervised learning of regression mixture models with unknown number of components

Faicel Chamroukhi

Aix Marseille Université, CNRS, ENSAM, LSIS, UMR 7296, 13397 Marseille, France
Université de Toulon, CNRS, LSIS, UMR 7296, 83957 La Garde, France
chamroukhi@univ-tln.fr

Abstract

Regression mixture models are widely studied in statistics, machine learning and data analysis. Fitting regression mixtures is challenging and is usually performed by maximum likelihood by using the expectation-maximization (EM) algorithm. However, it is well-known that the initialization is crucial for EM. If the initialization is inappropriately performed, the EM algorithm may lead to unsatisfactory results. The EM algorithm also requires the number of clusters to be given a priori; the problem of selecting the number of mixture components requires using model selection criteria to choose one from a set of pre-estimated candidate models. We propose a new fully unsupervised algorithm to learn regression mixture models with unknown number of components. The developed unsupervised learning approach consists in a penalized maximum likelihood estimation carried out by a robust expectation-maximization (EM) algorithm for fitting polynomial, spline and B-spline regressions mixtures. The proposed learning approach is fully unsupervised: 1) it simultaneously infers the model parameters and the optimal number of the regression mixture components from the data as the learning proceeds, rather than in a two-fold scheme as in standard model-based clustering using afterward model selection criteria, and 2) it does not require accurate initialization unlike the standard EM for regression mixtures. The developed approach is applied to curve clustering problems. Numerical experiments on simulated data show that the proposed robust EM algorithm performs well and provides accurate results in terms of robustness with regard initialization and retrieving the optimal partition with the actual number of clusters. An application to real data in the framework of functional data clustering, confirms the benefit of the proposed approach for practical applications.

1 Introduction

Mixture modeling (McLachlan and Peel., 2000) is one of the most popular and successful approaches in density estimation and cluster and discriminant analysis (McLachlan and Peel., 2000; Titterington et al., 1985; McLachlan and Basford, 1988; Banfield and Raftery, 1993; Fraley and Raftery, 1998, 2002, 2007). In this paper we focus on clustering using the finite mixture model (McLachlan and Peel., 2000), that is model-based clustering (Banfield and Raftery, 1993; McLachlan and Basford, 1988; Fraley and Raftery, 2002) which is one of the most popular and

successful approaches in cluster analysis. In the finite mixture approach for cluster analysis, the data probability density function is assumed to be a mixture density, each component density being associated with a cluster. The problem of clustering, therefore, becomes the one of estimating the parameters of the assumed mixture model (e.g, estimating the mean vector and the covariance matrix for each component density in the case of Gaussian mixture models). The mixture density estimation in this case is generally performed by maximizing the observed data log-likelihood. This can be achieved by the well-known expectation-maximization (EM) algorithm (Dempster et al., 1977; McLachlan and Krishnan, 1997) thanks to its good desirable properties of stability and reliable convergence. One of the main model-based clustering approaches is the one based on Gaussian mixture models and the EM algorithm or its extensions (McLachlan and Peel., 2000; McLachlan and Krishnan, 1997; Yang et al., 2012). It is in general concerned with multivariate data vectors. However, in many areas of application, the data are curves or functions. The analysis approaches in this case relate the Functional Data Analysis (FDA) framework (Ramsay and Silverman, 2005) which concerns the paradigm of data analysis for which the individuals are entire functions or curves rather than vectors of reduced dimensions. When the observations are curves or time series, the clustering can therefore be performed by using model-based curve clustering approaches, namely the regression mixture model including polynomial regression mixtures, spline and B-spline regression mixtures (Gaffney and Smyth, 1999; Gaffney, 2004; Chamroukhi, 2010). Modeling with regression mixtures is an important topic within the general family of mixture models. The regression mixture model (Quandt, 1972; Quandt and Ramsey, 1978; Veaux, 1989; Jones and McLachlan, 1992; Gaffney and Smyth, 1999; Viele and Tong, 2002; Faria and Soromenho, 2010; Chamroukhi, 2010; Young and Hunter, 2010; Hunter and Young, 2012) arises when we assume that the observed response $\mathbf{y}_i$ for the predictor variable $\mathbf{x}_i$ is generated from one of K possibly parametric regression functions $g(\mathbf{y}_i|\mathbf{x}_i;\boldsymbol{\theta}_k)$ of parameters $\boldsymbol{\theta}_k$ with prior probability π_k. This includes polynomial regression mixtures, spline regression mixtures and B-spline regression mixtures. The problem of fitting regression mixture models is a widely studied problem in statistics, machine learning and data analysis, in particular for cluster analysis. It is usually performed by maximum likelihood using the expectation-maximization (EM) algorithm (Dempster et al., 1977; Gaffney and Smyth, 1999; Jones and McLachlan, 1992; McLachlan and Krishnan, 1997). It is well-known that the initialization is crucial for EM. The EM algorithm also requires the number of clusters to be given a priori; the problem of selecting the number of components in this case requires using model selection criteria to choose one from a set of pre-estimated candidate models. Some solutions have been provided in Yang et al. (2012) and Figueiredo and Jain (2000) for model-based clustering with Gaussian mixture models for multivariate data. In this paper we focus on model-based clustering approaches using regression mixtures for curves or functional data, rather than vectorial data. We propose a new fully unsupervised learning algorithm to learn regression mixture models with unknown number of components. The developed unsupervised learning approach consists in a penalized maximum likelihood estimation carried out by a robust expectation-maximization (EM) algorithm for fitting polynomial, spline and B-spline regressions mixtures. The proposed learning approach is fully unsupervised. It simultaneously infers the model parameters and the optimal number of the regression mixture components from the data as the learning proceeds, rather than in a

two-fold scheme as in standard model-based clustering using afterward model selection criteria. Furthermore, it does not require accurate initialization unlike the standard EM for regression mixtures.

This paper is organized as follows. In section 2 we give a brief background on model-based clustering using regression mixtures. In section 3 we present the proposed approach and the robust EM algorithm which maximizes a penalized log-likelihood criterion for regression mixtures in model-based (curve) clustering, and derive the corresponding parameter updating formulas. Then, in section 4 we present experimental results on both simulated data and real data to evaluate the proposed approach. Finally, section 5 we discuss the proposal and provide concluding remarks and future directions.

2 Model-based clustering with regression mixtures

2.1 Model-based clustering

Model-based clustering is a probabilistic generative approach for clustering that is based on the finite mixture model formulation (Banfield and Raftery, 1993; McLachlan and Basford, 1988; Fraley and Raftery, 2002; McLachlan and Peel., 2000). It assumes that an observed data point $\mathbf{x}_i$ from a sample $(\mathbf{x}_1, \ldots, \mathbf{x}_n)$, which is in general a multidimensional vector in $\mathbb{R}^d$, is (independently) generated from one of K clusters with prior probability π_k. Each cluster being modeled by its specific density f_k with possibly a set of parameters $\boldsymbol{\theta}_k$. The data are thus modeled by a finite mixture model of K components f_k, which can be written as:

$$f(\mathbf{x}_i; \boldsymbol{\theta}) = \sum_{k=1}^{K} \pi_k f_k\big(\mathbf{x}_i; \boldsymbol{\theta}_k\big). \tag{1}$$

For example, in the case of a Gaussian mixture model, each of the density components $f_k\big(\mathbf{x}_i; \boldsymbol{\theta}_k\big)$ is a Gaussian density $\mathcal{N}(.; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ with parameters $\boldsymbol{\theta}_k = (\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ where $\boldsymbol{\mu}_k$ is the mean vector and $\boldsymbol{\Sigma}_k$ is the covariance matrix. The problem of clustering therefore becomes the one of estimating the mixture parameters $\boldsymbol{\theta}$. The parameter estimation can be performed by maximizing the observed-data log-likelihood:

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^{n} \log \sum_{k=1}^{K} \pi_k f_k\big(\mathbf{x}_i; \boldsymbol{\theta}_k\big) \tag{2}$$

by using the EM algorithm (Dempster et al., 1977; McLachlan and Krishnan, 1997).

The standard model-based clustering technique based on multivariate unconditional mixture (1) has been extended to the problem of clustering curves or functions, where the observed data are composed of both responses $\mathbf{y}_i$ and corresponding covariates $\mathbf{x}_i$. In this case, the mixture component densities $f_k\big(\mathbf{x}_i; \boldsymbol{\theta}_k\big)$ in (1) are replaced by conditional densities $f_k\big(\mathbf{y}_i|\mathbf{x}_i; \boldsymbol{\theta}_k\big)$, and a common way to model the conditional dependence in the observed data is to use regression functions. The resulting model is thus a regression mixture model, including polynomial regression mixture and (B-)spline regression mixtures, see for example (DeSarbo and Cron, 1988; Jones

and McLachlan, 1992; Gaffney, 2004). These approaches use the maximum likelihood estimation approach with the EM algorithm to estimate the model parameters. In the following section, we give an overview on the related model-based curve clustering approaches based on regression mixture models and the EM algorithm.

2.2 Regression mixtures for model-based clustering

The regression mixtures approaches assume that each observation $(\mathbf{x}_i, \mathbf{y}_i)$ is drawn from one of K clusters whose proportions (prior probabilities) are $(\pi_1, \ldots, \pi_K)$ and each cluster k is supposed to be a set of realizations from a polynomial regression model with coefficients $\boldsymbol{\beta}_k$ corrupted by a standard zero-mean Gaussian noise with a covariance-matrix $\sigma_k^2 \mathbf{I}_{m_i}$:

$$\mathbf{y}_i = \mathbf{X}_i \boldsymbol{\beta}_k + \boldsymbol{\epsilon}_{ik} \tag{3}$$

where $\mathbf{y}_i = (y_{i1}, \ldots, y_{im_i})^T$ is an $m_i \times 1$ vector, $\mathbf{X}_i$ is the $m_i \times (p+1)$ regression matrix (Vandermonde matrix) with rows

$$\boldsymbol{x}_j = (1, x_{ij}, x_{ij}^2 \ldots, x_{ij}^p), \tag{4}$$

p being the polynomial degree, $\boldsymbol{\beta}_k = (\beta_{k0}, \ldots, \beta_{kp})^T$ is the $(p+1) \times 1$ vector of regression coefficients for class k, $\boldsymbol{\epsilon}_{ik} \sim \mathcal{N}(\mathbf{0}, \sigma_k^2 \mathbf{I}_{m_i})$ is a multivariate standard zero-mean Gaussian vector with a covariance matrix $\sigma_k^2 \mathbf{I}_{m_i}$, representing the corresponding additive Gaussian noise and $\mathbf{I}_{m_i}$ denotes the $m_i \times m_i$ identity matrix. Thus, the conditional density of curves for cluster k is therefore given by the following functional normal density

$$f_k(\mathbf{y}_i | \mathbf{x}_i, z_i = k; \boldsymbol{\theta}_k) = \mathcal{N}(\mathbf{y}_i; \mathbf{X}_i \boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I}_{m_i}), \tag{5}$$

and the polynomial regression mixture is given by the following conditional mixture density:

$$f(\mathbf{y}_i | \mathbf{x}_i; \boldsymbol{\theta}) \quad = \quad \sum_{k=1}^{K} \pi_k \, \mathcal{N}(\mathbf{y}_i; \mathbf{X}_i \boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I}_{m_i}). \tag{6}$$

The model parameters are given by the parameter vector $\boldsymbol{\theta} = (\pi_1, \ldots, \pi_K, \boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_K)$ where the π_k's are the non-negative mixing proportions that sum to 1, $\boldsymbol{\theta}_k = (\boldsymbol{\beta}_k, \sigma_k^2)$, σ_k^2 represents the regression parameters and the noise variance for cluster k. The unknown parameter vector $\boldsymbol{\theta}$ is generally estimated by maximizing the observed-data log-likelihood given an *i.i.d.* dataset:

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^{n} \log \sum_{k=1}^{K} \pi_k \, \mathcal{N}(\mathbf{y}_i; \mathbf{X}_i \boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I}_{m_i}). \tag{7}$$

This is usually performed iteratively via the EM algorithm (Dempster et al., 1977; Gaffney and Smyth, 1999; Gaffney, 2004).

2.3 Standard EM for fitting regression mixtures

The EM algorithm is the standard tool to maximize the observed-data log-likelihood (7). This is performed iteratively by maximizing, at each EM step, the conditional expectation of the log-likelihood of $\boldsymbol{\theta}$ given the complete data, that is the complete-data log-likelihood, and a current parameter estimation of $\boldsymbol{\theta}$. The complete data in this case are composed of the observed data $\mathcal{D} = ((\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_n, \mathbf{y}_n))$ and the corresponding unknown cluster labels $\mathbf{z} = (z_1, \ldots, z_n)$ with $z_i \in \{1, \ldots, K\}$. Thus, the complete-data log-likelihood of $\boldsymbol{\theta}$ is given by:

$$\mathcal{L}_c(\boldsymbol{\theta}) = \sum_{i=1}^{n} \sum_{k=1}^{K} z_{ik} \log \left[\pi_k \mathcal{N}(\mathbf{y}_i; \mathbf{X}_i \boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I}_{m_i})\right] \tag{8}$$

where z_{ik} is an indicator binary-valued variable such that $z_{ik} = 1$ if $z_i = k$ (i.e., if $\mathbf{y}_i$ is generated by the polynomial regression component k) and $z_{ik} = 0$ otherwise.

The EM algorithm for regression mixtures starts with an initial parameter vector $\boldsymbol{\theta}^{(0)}$ and alternates between the two following steps until convergence:

E-step: This step computes the expected complete-data log-likelihood given the observed data $\mathcal{D}$ and a current estimation $\boldsymbol{\theta}^{(q)}$ of the parameter vector $\boldsymbol{\theta}$, q being the current iteration number:

$$\begin{aligned} Q(\boldsymbol{\theta}, \boldsymbol{\theta}^{(q)}) &= \mathbb{E}\big[\mathcal{L}_c(\boldsymbol{\theta}) | \mathcal{D}; \boldsymbol{\theta}^{(q)}\big] \\ &= \sum_{i=1}^{n} \sum_{k=1}^{K} \tau_{ik}^{(q)} \log \left[\pi_k \mathcal{N}(\mathbf{y}_i; \mathbf{X} \boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I}_{m_i})\right]. \end{aligned} \tag{9}$$

where

$$\begin{aligned} \tau_{ik}^{(q)} &= p(z_i = k | \mathbf{y}_i, \mathbf{x}_i; \boldsymbol{\theta}^{(q)}) \\ &= \frac{\pi_k^{(q)} \mathcal{N}\big(\mathbf{y}_i; \mathbf{X}_i \boldsymbol{\beta}_k^{(q)}, \sigma_k^{2(q)} \mathbf{I}_{m_i}\big)}{\sum_{h=1}^{K} \pi_h^{(q)} \mathcal{N}(\mathbf{y}_i; \mathbf{X}_i \boldsymbol{\beta}_h^{(q)}, \sigma_h^{2(q)} \mathbf{I}_{m_i})} \end{aligned} \tag{10}$$

is the posterior probability that the ith curve is generated from cluster k. This step therefore simply consists in computing the posterior cluster probabilities.

M-step: This step updates the model parameters and provides the parameter vector $\boldsymbol{\theta}^{(q+1)}$ by maximizing the Q-function (9) with respect to $\boldsymbol{\theta}$. The mixing proportions are updated by maximizing $\sum_{i=1}^{n} \sum_{k=1}^{K} \tau_{ik}^{(q)} \log \pi_k$ with respect to $(\pi_1, \ldots, \pi_K)$ subject to the constraint $\sum_{k=1}^{K} \pi_k = 1$ using Lagrange multipliers, which gives the following updates (McLachlan and Krishnan, 1997; McLachlan and Peel., 2000):

$$\pi_k^{(q+1)} = \frac{1}{n} \sum_{i=1}^{n} \tau_{ik}^{(q)} \quad (k = 1, \ldots, K). \tag{11}$$

The maximization of each of the K functions $\sum_{i=1}^{n} \tau_{ik}^{(q)} \log \mathcal{N}(\mathbf{y}_i; \mathbf{X}\boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I}_{m_i})$ in (9) consists in solving a weighted least-squares problem. The solution of this problem is straightforward and the updating equations are equivalent to the well-known weighted least-squares solutions (see for example (Gaffney, 2004; Chamroukhi, 2010)):

$$\boldsymbol{\beta}_k^{(q+1)} = \Big[\sum_{i=1}^{n} \tau_{ik}^{(q)} \mathbf{X}_i^T \mathbf{X}_i\Big]^{-1} \sum_{i=1}^{n} \tau_{ik}^{(q)} \mathbf{X}_i^T \mathbf{y}_i \tag{12}$$

$$\sigma_k^{2(q+1)} = \frac{1}{\sum_{i=1}^{n} \tau_{ik}^{(q)}} \sum_{i=1}^{n} \tau_{ik}^{(q)} \parallel \mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}_k \parallel^2 . \tag{13}$$

From a clustering prospective, the estimated mixture components, are interpreted as K clusters, where each cluster is associated to a mixture component. Thus, once the model parameters have been estimated, the posterior probability that an individual was generated by the kth cluster (or component) given by (10) represent a fuzzy partition of the data into K clusters. The posterior cluster probabilities indeed represent a fuzzy cluster memberships. The represent the a posteriori uncertainty given the observed data and the estimated model about which cluster k each observed data $(\mathbf{x}_i, \mathbf{y}_i)$ is originated from. A hard partition can then be obtained by assigning each observed curve to the cluster with the highest posterior probability:

$$\hat{z}_i = \arg \max_{k=1}^{K} \hat{\tau}_{ik}. \tag{14}$$

The pseudo code 1 summarizes the standard EM algorithm for regression mixtures.

2.4 Spline and B-spline regression mixtures

The (B-)spline regression mixture is an extension of the previously described polynomial regression mixture to a semiparametric modeling by relying on splines or B-splines.

2.4.1 Spline and B-spline regression

Splines (Deboor, 1978) are widely used for function approximation based on constrained piecewise polynomials. Let $\boldsymbol{\xi} = \xi_0 < \xi_1, ..., < \xi_L < \xi_{L+1}$ be ordered knots, including L internal knots, ξ_0 and ξ_{L+1} being the two boundary knots. An order-M spline with knots $\boldsymbol{\xi}$ is a piecewise-polynomial of degree $p = M - 1$ with continuous derivatives at the interior knots up to order $M - 2$. For example an order-2 spline is a continuous piecewise linear function. The spline regression function can be written as:

$$y_{ij} = \sum_{\ell=0}^{p} \beta_\ell x_{ij}^\ell + \sum_{\ell=1}^{L} \beta_{\ell+p} (x_{ij} - \xi_\ell)_+^p + \epsilon_{ij} \tag{15}$$

where $(x_{ij} - \xi_\ell)_+ = x_{ij} - \xi_\ell$ if $x_{ij} \geq \xi_k$ and $(x_{ij} - \xi_\ell)_+ = 0$ otherwise, $\boldsymbol{\beta} = (\beta_0, \ldots, \beta_{L+p})^T$ is the vector of spline coefficients in $\mathbb{R}^{L+M}$ and ϵ is and additive Gaussian noise. This spline regression

Algorithm 1 Pseudo code of the standard EM algorithm for regression mixtures.

Inputs: Set of curves $\mathcal{D} = ((\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_n, \mathbf{y}_n))$, number of clusters K and polynomial degree p

fix a threshold $\epsilon > 0$; set $q \leftarrow 0$ (iteration)
Initialization: $\boldsymbol{\theta}^{(0)} = (\pi_1^{(0)}, \ldots, \pi_K^{(0)}, \boldsymbol{\theta}_1^{(0)}, \ldots, \boldsymbol{\theta}_K^{(0)})$:
Initialize randomly the partition or by running K-means and initialize the π_k's
Fit a regression model with parameters $\boldsymbol{\theta}_k^{(0)} = (\boldsymbol{\beta}_k^{(0)}, \sigma_k^{2(0)})$ to each cluster k
while increment in log-likelihood $> \epsilon$ **do**
  `//E-step:`
  for $k = 1, \ldots, K$ **do**
    Compute $\tau_{ik}^{(q)}$ for $i = 1, \ldots, n$ using Equation (10)
  end for
  `//M-step:`
  for $k = 1, \ldots, K$ **do**
    Compute $\pi_k^{(q+1)}$ using Equation (11)
    Compute $\boldsymbol{\beta}_k^{(q+1)}$ using Equation (13)
    Compute $\sigma_k^{2(q+1)}$ using Equation (13)
  end for
  $q \leftarrow q + 1$
end while
Outputs: $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}^{(q)}$, $\hat{\tau}_{ik} = \tau_{ik}^{(q)}$

model can be written in a vectorial form as

$$\mathbf{y}_i = \mathbf{S}_i \boldsymbol{\beta} + \boldsymbol{\epsilon}_i \tag{16}$$

where $\mathbf{S}_i$ is the $m_i \times (L + M)$ spline regression matrix with rows

$$\mathbf{s}_j = (1, x_{ij}, x_{ij}^2 \ldots, x_{ij}^p, (x_{ij} - \xi_1)_+^p, \ldots, (x_{ij} - \xi_L)_+^p). \tag{17}$$

For splines, the columns of the regression matrix $\mathbf{X}$ tend to be highly correlated since each column is a transformed version of x. This collinearity may result in a nearly singular matrix and imprecision in the spline fit (Ruppert and Carroll, 2003). B-splines allows for efficient computations thanks to the block matrix form of the regression matrix. An order-M B-spline function is defined as a sum of linear combination of specific basis functions $B_{\ell,M}$ as:

$$y_{ij} = \sum_{\ell=1}^{L+M} \beta_\ell B_{\ell,M}(x_{ij}), \; x_{ij} \in [\zeta_\ell, \zeta_{\ell+M}] \tag{18}$$

where each Mth order B-spline $B_{\ell,M}$ is a piecewise polynomial of degree $p = M - 1$ that has finite support over $[\zeta_l, \zeta_{\ell+M}]$ and is zero everywhere else. Each of the basis functions $B_{\ell,M}(x_{ij})$ can be computed as in Appendix A (Hastie et al., 2010). For the B-spline regression model, each

row of the $m_i \times (L+M)$ B-spline regression matrix $\mathbf{B}_i$ for the ith curve is thus given by:

$$\mathbf{b}_j = (B_{1,M}(x_{ij}),\ B_{2,M}(x_{ij}), \ldots,\ B_{L+M,M}(x_{ij})). \tag{19}$$

2.4.2 Spline and B-spline regression mixtures and the EM algorithm

The spline regression mixture (respectively the B-spline regression mixture) is similarly derived as the polynomial regression mixture described previously. The mixture density in this case is given by Eq. (6) where the regression matrix $\mathbf{X}_i$ is replaced by the spline regression matrix $\mathbf{S}_i$ (respectively the B-slipne regression matrix $\mathbf{B}_i$).

The parameter estimation procedure for the (B-)spline regression mixture is the same as the one used for polynomial regression mixture, that is maximum likelihood estimation via the EM algorithm. The E- and M-steps are still the same, as well as as the initialization procedure and the convergence conditions.

However, it can be noticed that, the standard EM algorithm for all theses regression mixture models is sensitive to initialization. It might yield poor estimations if the mixture parameters are not initialized properly. The EM initialization in general can be performed from a randomly chosen partition of the data or by computing a partition from another clustering algorithm such as K-means, Classification EM (Celeux and Diebolt, 1985), Stochastic EM (Celeux and Govaert, 1992), etc or by initializing EM with a few number of steps of EM itself. Several works have been proposed in the literature in order to overcome this drawback and making the EM algorithm for Gaussian mixtures robust with regard initialization (Biernacki et al., 2003; Reddy et al., 2008; Yang et al., 2012). Further details about choosing starting values for the EM algorithm for Gaussian mixtures can be found for example in (Biernacki et al., 2003). In addition, the EM algorithm requires the number of mixture components (clusters) to be known. While the number of cluster can be chosen by some model selection criteria such as BIC (Schwarz, 1978), AIC (Akaike, 1974) or ICL (Biernacki et al., 2000), this requires performing an afterward model selection procedure. Some authors have considered this issue in order to estimate the unknown number of mixture components in Gaussian mixture models, for example as in (Richardson and Green, 1997; Yang et al., 2012).

In general, theses two issues have been considered each separately. Among the approaches considering both the problem of robustness with regard to initial values and automatically estimating the number of mixture components in the same algorithm, one can cite the EM algorithm proposed in Figueiredo and Jain (2000). The authors proposed an EM algorithm that is capable of selecting the number of components and it attempts to be not sensitive with regard to initial values. Indeed, the algorithm developed by Figueiredo and Jain (2000) optimizes a particular criterion called the minimum message length (MML), which is a penalized negative log-likelihood rather than the observed data log-likelihood. The penalization term allows to control the model complexity. It starts by fitting a mixture model with large number of clusters and discards illegitimate clusters as the learning proceeds. The degree of legitimate of each cluster is measured through the penalization term which includes the mixing proportions to know if the cluster is small or not to be discarded, and therefore to reduce the number of clusters. More recently, in

Yang et al. (2012), the authors developed a robust EM algorithm for model-based clustering of multivariate data using Gaussian mixture models that simultaneously addresses the problem of initialization and estimating the number of mixture components. This algorithm overcome some initialization drawback of the EM algorithm proposed in Figueiredo and Jain (2000), which is still having an initialization problem. As shown in Yang et al. (2012), this problem can become more serious especially for a dataset with a large number of clusters. However, these presented model-based clustering approaches, namely in Yang et al. (2012) and Figueiredo and Jain (2000), are concerned with vectorial data where the observations are assumed to be vectors of reduced dimension. When the data are rather curves or functions, they are not adapted. Indeed, when the data are functional where the individuals are presented as curves or surfaces rather than vectors, they are in general very structured. Relying on standard multivariate mixture analysis may therefore lead to unsatisfactory results in terms of modeling and classification accuracy (Chamroukhi et al., 2009, 2010, 2011). However, addressing the problem from a functional data analysis prospective, that is formulating "functional" mixture models, allows to overcome these limitations, e.g., (Chamroukhi et al., 2009, 2010, 2011; Gaffney, 2004). In this case of model-based functional data clustering, one can rely on the regression mixture approaches (Gaffney and Smyth, 1999; Gaffney, 2004; Chamroukhi, 2010) or generative hidden process regression (Chamroukhi et al., 2010; Chamroukhi, 2010) which are adapted for curves with regime changes. In this paper, we attempt to overcome the limitations of the EM algorithm in the case of regression mixtures and model-based curve clustering by proposing an EM algorithm which is robust with regard initialization and automatically estimates the optimal number of clusters as the learning proceeds. We consider the polynomial, spline and B-spline regression mixtures.

In the next section we derive our robust EM algorithm for fitting regression mixtures including polynomial, spline and B-spline regression mixtures.

3 Penalized maximum likelihood via a robust EM algorithm for fitting regression mixtures

In this section we present the proposed EM algorithm for model-based curve clustering using regression mixtures. The present work is in the same spirit of the EM algorithm presented in Yang et al. (2012) but by extending the idea to the case of functional data (curve) clustering rather than multivariate data clustering. It is therefore concernned with regression mixture models rather than multivariate GMMs. Indeed, the data here are assumed to be curves rather than vectors of reduced dimensions. This leads to fitting regression mixture models (including splines or B-splines), rather than fitting standard Gaussian mixtures. We start by describing the maximized objective function and then we derive the proposed EM algorithm to estimate the regression mixture model parameters.

3.1 Penalized maximum likelihood estimation

For estimating the regression mixture model (6), we attempt to maximize a penalized log-likelihood function rather than the standard observed-data log-likelihood (7). This criterion

consists in penalizing the observed-data log-likelihood (7) by a term accounting for the model complexity. As the model complexity is related to in particular the number of clusters and therefore the structure of the hidden variables z_i, we chose to use the entropy of the hidden variable z_i as penalty, $z_i$being the class label of the ith curve. The penalized log-likelihood criterion is therefore derived as follows. The discrete-valued variable $z_i \in \{1, \ldots, K\}$ with probability distribution $p(z_i)$ represents the hidden class label of the ith observed curve. The (differential) entropy of this variable is defined by

$$\begin{aligned} H(z_i) &= -\mathbb{E}[\log p(z_i)] = -\sum_{k=1}^{K} p(z_i = k) \log p(z_i = k) \\ &= -\sum_{k=1}^{K} \pi_k \log \pi_k. \end{aligned} \tag{20}$$

By assuming that the variables $\mathbf{z} = (z_1, \ldots, z_n)$ are independent and identically distributed (i.i.d), which is in general the assumption in clustering using mixtures where the cluster labels are assumed to be distributed according to a Multinomial distribution, the whole entropy for $\mathbf{z}$ is therefore additive and we have

$$H(\mathbf{z}) = -\sum_{i=1}^{n} \sum_{k=1}^{K} \pi_k \log \pi_k. \tag{21}$$

The penalized log-likelihood function we propose to maximize is thus constructed by penalizing the observed data log-likelihood (7) by the entropy term (21), that is

$$\mathcal{J}(\lambda, \boldsymbol{\theta}) = \mathcal{L}(\boldsymbol{\theta}) - \lambda H(\mathbf{z}), \quad \lambda \geq 0 \tag{22}$$

which leads to the following penalized log-likelihood criterion:

$$\mathcal{J}(\lambda, \boldsymbol{\theta}) = \sum_{i=1}^{n} \log \sum_{k=1}^{K} \pi_k \mathcal{N}(\mathbf{y}_i; \mathbf{X}_i \boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I}_{m_i}) + \lambda n \sum_{k=1}^{K} \pi_k \log \pi_k \tag{23}$$

where $\mathcal{L}(\boldsymbol{\theta})$ is the observed-data log-likelihood maximized by the standard EM algorithm for regression mixtures (see Equation (7)) and $\lambda \geq 0$ is a parameter of control that controls the complexity of the fitted model. This penalized log-likelihood function (23) we attempt to optimize allows to control the complexity of the model fit through the roughness penalty $\lambda n \sum_{k=1}^{K} \pi_k \log \pi_k$. The entropy term $-\sum_{i=1}^{n} \sum_{k=1}^{K} \pi_k \log \pi_k$ in the penalty measures the complexity of a fitted model for K clusters. When the entropy is large, the fitted model is rougher, and when it is small, the fitted model is smoother. The non-negative smoothing parameter λ is for establishing a trade-off between closeness of fit to the data and a smooth fit. As λ decreases, the fitted model tends to be less complex, and we get a smooth fit. However, when λ increases, the result is a rough fit. We discuss in the next section how to set this regularization coefficient in an adaptive way. The next section shows how the penalized observed-data log-likelihood $\mathcal{J}(\lambda, \boldsymbol{\theta})$ is maximized w.r.t the model parameters $\boldsymbol{\theta}$ by a robust EM algorithm for regression mixtures and model-based curve

clustering.

3.2 Robust EM algorithm for model-based curve clustering suing regression mixtures

Given an i.i.d training dataset of n curves $\mathcal{D} = ((\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_n, \mathbf{y}_n))$, the penalized log-likelihood (23) is iteratively maximized by using the following robust EM algorithm for model-based curve clustering. Before giving the EM steps, we give the penalized complete-data log-likelihood, on which the EM formulation is based. The complete-data log-likelihood, in this penalized case, is given by

$$\mathcal{J}_c(\lambda, \boldsymbol{\theta}) = \sum_{i=1}^{n}\sum_{k=1}^{K} z_{ik} \log\left[\pi_k \mathcal{N}(\mathbf{y}_i; \mathbf{X}_i\boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I}_{m_i})\right] + \lambda n \sum_{k=1}^{K} \pi_k \log \pi_k. \tag{24}$$

After starting with an initial solution (see section 3.3 for the initialization strategy and stopping rule), the proposed algorithm alternates between the two following steps until convergence.

3.2.1 E-step

This step computes the expectation of complete-data log-likelihood (24) over the hidden variables z_i, given the observations $\mathcal{D}$ and the current parameter estimation $\boldsymbol{\theta}^{(q)}$, q being the current iteration number:

$$\begin{aligned} Q(\lambda, \boldsymbol{\theta}; \boldsymbol{\theta}^{(q)}) &= \mathbb{E}\left[\mathcal{J}_c(\lambda, \boldsymbol{\theta}) | \mathcal{D}; \boldsymbol{\theta}^{(q)}\right] \\ &= \sum_{i=1}^{n}\sum_{k=1}^{K} \mathbb{E}\left[z_{ik} | \mathcal{D}; \boldsymbol{\theta}^{(q)}\right] \log\left[\pi_k \mathcal{N}(\mathbf{y}_i; \mathbf{X}_i\boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I}_{m_i})\right] + \lambda n \sum_{k=1}^{K} \pi_k \log \pi_k \\ &= \sum_{i=1}^{n}\sum_{k=1}^{K} \tau_{ik}^{(q)} \log \mathcal{N}(\mathbf{y}_i; \mathbf{X}_i\boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I}_{m_i}) + \sum_{i=1}^{n}\sum_{k=1}^{K} \tau_{ik}^{(q)} \log \pi_k + \lambda n \sum_{k=1}^{K} \pi_k \log \pi_k \end{aligned} \tag{25}$$

where

$$\tau_{ik}^{(q)} = p(z_i = k | \mathbf{y}_i, \mathbf{x}_i; \boldsymbol{\theta}^{(q)}) = \frac{\pi_k^{(q)} \mathcal{N}\left(\mathbf{y}_i; \mathbf{X}_i\boldsymbol{\beta}_k^{T(q)}, \sigma_k^{2(q)} \mathbf{I}_{m_i}\right)}{\sum_{h=1}^{K} \pi_h^{(q)} \mathcal{N}(\mathbf{y}_i; \mathbf{X}_i\boldsymbol{\beta}_h^{(q)}, \sigma_h^{2(q)} \mathbf{I}_{m_i})} \tag{26}$$

is the posterior probability that the curve $(\mathbf{x}_i, \mathbf{y}_i)$ is generated by the cluster k. This step therefore only requires the computation of the posterior cluster probabilities $\tau_{ik}^{(q)}$ $(i = 1, \ldots, n)$ for each of the K clusters.

3.2.2 M-step

This step updates the value of the parameter vector $\boldsymbol{\theta}$ by maximizing the Q-function (25) with respect to $\boldsymbol{\theta}$, that is: $\boldsymbol{\theta}^{(q+1)} = \arg\max_{\boldsymbol{\theta}} Q(\lambda, \boldsymbol{\theta}; \boldsymbol{\theta}^{(q)})$. It can be shown that this maximization can be performed by separate maximizations w.r.t the mixing proportions $(\pi_1, \ldots, \pi_K)$ subject to the constraint $\sum_{k=1}^{K} \pi_k = 1$, and w.r.t the regression parameters $\{\boldsymbol{\beta}_k, \sigma_k^2\}$. The mixing proportions

updates are obtained by maximizing the function

$$Q_\pi(\lambda;\boldsymbol{\theta}^{(q)}) = \sum_{i=1}^{n}\sum_{k=1}^{K}\tau_{ik}^{(q)}\log\pi_k + \lambda n\sum_{k=1}^{K}\pi_k\log\pi_k$$

w.r.t the mixing proportions $(\pi_1,\ldots,\pi_K)$ subject to the constraint $\sum_{k=1}^{K}\pi_k = 1$. This can be solved using Lagrange multipliers (see Appendix B) and the obtained updating formula is given by:

$$\pi_k^{(q+1)} = \frac{1}{n}\sum_{i=1}^{n}\tau_{ik}^{(q)} + \lambda\pi_k^{(q)}\left(\log\pi_k^{(q)} - \sum_{h=1}^{K}\pi_h^{(q)}\log\pi_h^{(q)}\right). \tag{27}$$

We notice here that in the update of the mixing proportions (27) the update is close to the standard EM update $\left(\frac{1}{n}\sum_{i=1}^{n}\tau_{ik}^{(q)}\right.$ see Eq. (11)$\left.\right)$ for very small value of λ. However, for a large value of λ, the penalization term will play its role in order to make clusters competitive and thus allows for discarding illegitimate clusters and enhancing actual clusters. Indeed, in the updating formula (27), we can see that for cluster k if

$$\left(\log\pi_k^{(q)} - \sum_{h=1}^{K}\pi_h^{(q)}\log\pi_h^{(q)}\right) > 0 \tag{28}$$

that is, for the (logarithm of the) current proportion $\log\pi_k^{(q)}$, the entropy of the hidden variables is decreasing, and therefore the model complexity tends to be stable, the cluster k has therefore to be enhanced. This explicitly results in the fact that the update of the kth mixing proportion $\pi_k^{(q+1)}$ in (27) will increase. On the other hand, if (28) is less than 0, the cluster proportion will therefore decrease as is not very informative in the sense of the entropy.

Finally, the penalization coefficient λ has to be set as described previously in such a way to be large for enhancing competition when the proportions are not increasing enough from one iteration to another. In this case, the robust algorithm plays its role for estimating the number of clusters (which is decreasing in this case by discarding small illegitimate clusters). We note that here a cluster k can be discarded if its proportion is less than $\frac{1}{n}$, that is $\pi_k^{(q)} < \frac{1}{n}$. On the other hand, λ has to become small when the proportions are sufficiently increasing as the learning proceeds and the partition can therefore be considered as stable. In this case, the robust EM algorithm tends to have the same behavior as the stand EM described in section 2.3. The regularization coefficient λ can be set in $[0,1]$ to prevent very large values. Furthermore, it can be set in an adaptive way similarly as described in Yang et al. (2012) for multivariate Gaussian mixtures according to the following adaptive formula to adapt it as the learning proceeds:

$$\lambda^{(q+1)} = \min\left\{\frac{\sum_{k=1}^{K}\exp\left(\eta n|\pi_k^{(q+1)} - \pi_k^{(q)}|\right)}{K}, \frac{1-\max_{k=1}^{K}\left(\frac{\sum_{i=1}^{n}\tau_{ik}^{(q)}}{n}\right)}{-\pi_k^{(q)}\sum_{k=1}^{K}\pi_k^{(q)}\log\pi_k^{(q)}}\right\} \tag{29}$$

where η can be set as $\min(1, 0.5^{\lfloor\frac{m}{2}-1\rfloor})$, m being the number of observations per curve and $\lfloor x\rfloor$ denotes the largest integer that is no more than x.

The maximization w.r.t the regression parameters however consists in separately maximizing

for each class k the function

$$\begin{aligned} Q_{\boldsymbol{\theta}_k}(\lambda, \boldsymbol{\beta}_k, \sigma_k^2; \boldsymbol{\theta}^{(q)}) &= \sum_{i=1}^{n} \tau_{ik}^{(q)} \log \mathcal{N}(\mathbf{y}_i; \mathbf{X}_i\boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I}_{m_i}) \\ &= \sum_{i=1}^{n} \tau_{ik}^{(q)} \left[-\frac{m}{2} \log 2\pi - \frac{m}{2} \log \sigma_k^2 - \frac{1}{2\sigma_k^2} \parallel \mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta}_k \parallel^2 \right] \end{aligned}$$

w.r.t $(\boldsymbol{\beta}_k, \sigma_k^2)$. This maximization w.r.t the regression parameters yields to analytic solutions of weighted least-squares problems where the weights are the posterior cluster probabilities $\tau_{ik}^{(q)}$. The updating formula are given by:

$$\boldsymbol{\beta}_k^{(q+1)} = \left[\sum_{i=1}^{n} \tau_{ik}^{(q)} \mathbf{X}_i^T \mathbf{X}_i \right]^{-1} \sum_{i=1}^{n} \tau_{ik}^{(q)} \mathbf{X}_i^T \mathbf{y}_i \tag{30}$$

$$\sigma_k^{2(q+1)} = \frac{1}{m \sum_{i=1}^{n} \tau_{ik}^{(q)}} \sum_{i=1}^{n} \tau_{ik}^{(q)} \parallel \mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta}_k \parallel^2 \tag{31}$$

where the posterior cluster probabilities $\tau_{ik}^{(q)}$ given by (26) are computed using the mixing proportions derived in (27). Then, once the model parameters have been estimated, a fuzzy partition of the data into K clusters is then represented by the posterior cluster probabilities τ_{ik}. A hard partition can then be computed according to the MAP principle by maximizing the posterior cluster probabilities, that is

$$\hat{z}_i = \arg\max_k \tau_{ik} \tag{32}$$

where $\hat{z}_i$ denotes the estimated cluster label of the ith observation.

3.3 Initialization strategy and stopping rule

The initial number of clusters is $K^{(0)} = n$, n being the total number of curves and the initial mixing proportions are $\pi_k^{(0)} = \frac{1}{K^{(0)}}$, $(k = 1, \ldots, K^{(0)})$. Then, to initialize the regression parameters $\boldsymbol{\beta}_k$ and the noise variances $\sigma_k^{2(0)}$ $(k = 1, \ldots, K^{(0)})$, we fitted a polynomial regression models on each curve k, $(k = 1, \ldots, K^{(0)})$ and the initial values of the regression parameters are therefore given by $\boldsymbol{\beta}_k^{(0)} = \left(\mathbf{X}_k^T \mathbf{X}_k\right)^{-1} \mathbf{X}_k \mathbf{y}_k$ and the noise variance can be deduced as $\sigma_k^{2(0)} = \frac{1}{m} \parallel \mathbf{y}_k - \mathbf{X}_k\boldsymbol{\beta}_k^{(0)} \parallel^2$. However, to avoid singularities at the starting point, we set $\sigma_k^{2(0)}$ as a middle value in the following sorted range $\parallel \mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta}_k^{(0)} \parallel^2$ for $i = 1, \ldots, n$.

The proposed EM algorithm is stopped when the maximum variation of the estimated regression parameters between two iterations $\max_{1 \leq k \leq K^{(q)}} \parallel \boldsymbol{\beta}_k^{(q+1)} - \boldsymbol{\beta}_k^{(q)} \parallel$ was less than a fixed threshold ϵ (e.g., 10^{-6}). The pseudo code 2 summarizes the proposed robust EM algorithm for model-based curve clustering using regression mixtures.

Algorithm 2 Pseudo code of the proposed Robust EM algorithm for polynomial regression mixtures.

Inputs: Set of curves $\mathcal{D} = ((\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_n, \mathbf{y}_n))$ and polynomial degree p

$\epsilon \leftarrow 10^{-6}$; $q \leftarrow 0$; $converge \leftarrow 0$
`//Initialization:`
$\boldsymbol{\theta}^{(0)} = (\pi_1^{(0)}, \ldots, \pi_K^{(0)}, \boldsymbol{\theta}_1^{(0)}, \ldots, \boldsymbol{\theta}_K^{(0)})$; $K^{(0)} = n$
for $k = 1, \ldots, K^{(q)}$ **do**
  Compute $\tau_{ik}^{(q)}$ for $i = 1, \ldots, n$ using Equation (26)
  Compute $\pi_k^{(q)}$ using Equation (27)
end for
`//EM loop:`
while (! $converge$) **do**
  for $k = 1, \ldots, K^{(q)}$ **do**
    Compute $\tau_{ik}^{(q)}$ for $i = 1, \ldots, n$ using Equation (26)
  end for
  for $k = 1, \ldots, K^{(q)}$ **do**
    Compute $\pi_k^{(q+1)}$ using Equation (27)
  end for
  Compute $\lambda^{(q+1)}$ using Equation (29) ; Discard illegitimate clusters with small proportions $\pi_k^{(q)} < \frac{1}{n}$; Set $K^{(q+1)} = K^{(q)} - \#\{\pi_k^{(q)} | \pi_k^{(q)} < \frac{1}{n}\}$; normalize $\tau_{ik}^{(q+1)}$ and $\pi_k^{(q+1)}$ so that their columns sum to one
  for $k = 1, \ldots, K^{(q)}$ **do**
    Compute $\boldsymbol{\beta}_k^{(q+1)}$ using Equation (30)
    Compute $\sigma_k^{2(q+1)}$ using Equation (31)
  end for
  if $\max_{1 \leq k \leq K^{(q)}} \| \boldsymbol{\beta}_k^{(q+1)} - \boldsymbol{\beta}_k^{(q)} \| < \epsilon$ **then**
    $converge = 1$
  end if
  $q \leftarrow q + 1$
end while

Outputs: $K = K^{(q)}$, $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}^{(q)}$, $\hat{\tau}_{ik} = \tau_{ik}^{(q)}$

3.4 Choosing the order of regression and spline knots number and locations

For a general use of the proposed algorithm for the polynomial regression mixture, the order of regression can be chosen by cross-validation techniques as in (Gaffney, 2004). In our experiments, we report the results corresponding to the for which th polynomial regression mixture provides the best fit. However, in some situations, the PRM model may be too simple to capture the full structure of the data, in particular for curves with high non-linearity of with regime changes, even if they can be seen as providing a useful first-order approximation of the data structure. The (B)-spline regression models in this case are more adapted. For these models, on may need to

choose the spline order and the number of knots and their locations. For the order of regression in (B)-splines, we notice that, in practice, the most widely used orders are $M = 1,2$ and 4 (Hastie et al., 2010). For smooth functions approximation, cubic (B)-splines (of order 4), which correspond to a (B)-spline of order 4 and thus with twice continuous derivatives, are sufficient to approximate smooth functions. When the data present irregularity, such as a kind of piecewise non continuous functions, a linear spline (of order 2) should be more adapted. This was namely used for the satellite dataset. The order 1 can be chosen for piecewise constant data. Concerning the choice of the number of knots and their locations, a common choice is to place a number of knots uniformly spaced across the range of x. In general more knots are needed for functions with regime changes. One can also use automatic techniques for the selection of the number of knots and their locations as reported in Gaffney (2004). For example, this can be performed by using cross validation as in Ruppert and Carroll (2003). In Kooperberg and Stone (1991), the knots are placed at selected order statistics of the sample data and the number of knots is determined either by a simple rule or by minimizing a variant of AIC. The general goal is to use a sufficient number of knots to fit the data while and at the same time to avoid over-fitting and to not make the computation excessive. The current algorithm can be easily extended to handle this type of automatic selection of spline knots placement, but as the unsupervised clustering problem itself requires much attention and is difficult, it is wise to fix the number and location of knots. In this paper, we will use knot sequences which are uniformly spaced across the range of x. The studied problems are not not very sensitive to the number and location of knots. Few number of equispaced knots (less than ten for the data studied in this paper) are sufficient to fit the data

4 Experimental study

This section is dedicated to the evaluation of the proposed approach on simulated data and real-world data. The algorithms have been implemented on Matlab and the esperiments were performed on a personal laptop. The developed Matlab codes are available upon request from the author. We evaluate the proposed EM algorithm for the three regression mixtures models: the polynomial regression mixture, spline regression mixture and B-spline regression mixture, respectively abbreviated as PRM, SRM and bSRM. The evaluation is performed in terms of estimating the actual partition by considering the estimated number of clusters and the clustering accuracy (misclassification error). We first consider simulated data and the waveforms benchmark. Then, we consider three real-world data sets issued covering three different application area: phoneme recognition, clustering gene time course data for bio-informatics and clustering satellite data.

4.1 Simulation study

The simulated data are simulated arbitrary curves and a the well-known benchmark of waveform data of Breiman et al. (1984a).

4.1.1 Experiments on simulated curves

In this experiment we consider non-linear arbitrary curves simulated as follows. The curves are simulated from situation representing a three-class problem and consisted of $n = 100$ arbitrary non-linear curves, each curve is composed of $m = 50$ observations. For the ith curve ($i = 1, \ldots, n$), the jth observation ($j = 1, \ldots, m$) is generated as follows:

- $y_{ij} = 0.8 + 0.5 \exp(-1.5\, x) \sin(1.3\pi\, x) + \sigma_1\, \epsilon_{ij}$;
- $y_{ij} = 0.5 + 0.8 \exp(-x) \sin(0.9\pi\, x) + \sigma_2\, \epsilon_{ij}$;
- $y_{ij} = 1 + 0.5 \exp(-x) \sin(1.2\pi\, x) + \sigma_3\, \epsilon_{ij}$;

with $\sigma_1 = \sigma_2 = \sigma_3 = 0.1$. The classes have respectively proportions $\pi_1 = 0.4, \pi_2 = 0.3, \pi_3 = 0.3$. Figure 1 shows a sample of the simulated data.

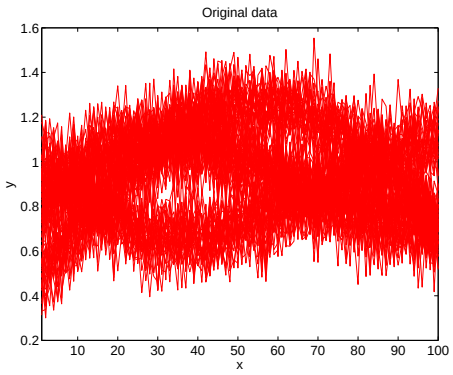

Figure 1: Simulated arbitrary nonlinear curves from a three-class problem.

Figure 2 shows the clustering results obtained by a polynomial regression mixture model, a cubic spline and B-spline regression mixtures. Each cluster is shown with its estimated mean function and the confidence interval of $\pm 2\hat{\sigma}$. It can be observed that the actual number of clusters (three) is correctly estimated by the proposed algorithm for the three models. for this simulated sample, the actual partition is also correctly recovered for the three models and all the observations are correctly classified.

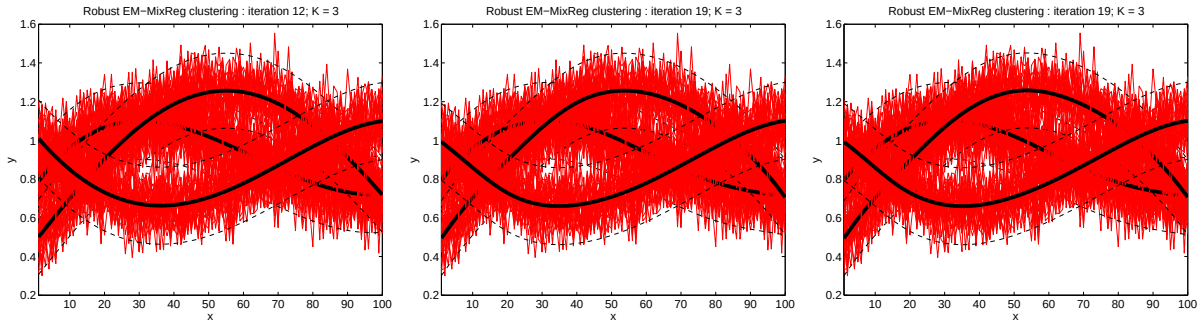
Robust EM−MixReg clustering : iteration 12; K = 3
Robust EM−MixReg clustering : iteration 19; K = 3
Robust EM−MixReg clustering : iteration 19; K = 3
y
x

Figure 2: Clustering results obtained with a polynomial regression mixture (PRM) (left) with $p = 4$, a cubic spline regression mixture (SRM) (middle) and a cubic B-spline regression mixture (bSRM) (right) with 4 knots.

Figure 3 shows the actual and estimated cluster mean functions for each regression mixture model. We ccan observed that the results are quasi-identical and the real mean functions are estimated accurately.

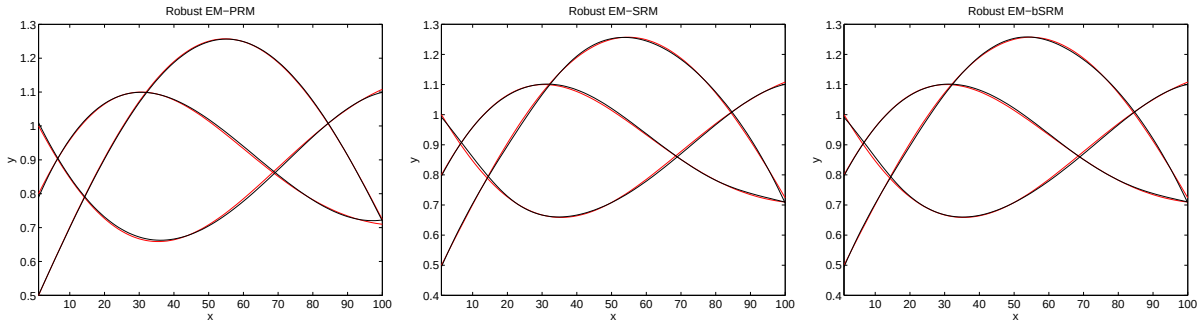

Figure 3: Actual (red) and estimated cluster mean functions (black) for the polynomial regression PRM (left), the SRM (middle) and the bSRM (right).

Table 1 gives numerical results for this data set. We include the approximation error which is computed as the normalized squared error between the true cluster mean functions and the estimated ones. The data are correctly classified by the three models and the actual number of clusters is correctly estimated. We can also see that the errors in terms of approximation of the true cluster mean functions are very low, as well as the ones of the true cluster variance estimation.

	K	misc. error rate	σ_1	σ_2	σ_3	π_1	π_2	π_3	Approx. Error
actual	3	–	1	1	1	0.4	0.3	0.3	–
EM-PRM	3	0%	0.0988	0.0969	0.0995	0.4	0.3	0.3	$4.4979e-05$
EM-SRM	3	0%	0.0987	0.0968	0.0994	0.4	0.3	0.3	$6.8076e-05$
EM-bSRM	3	0%	0.0987	0.0968	0.0994	0.4	0.3	0.3	$6.4905e-05$

Table 1: Clustering results for the simulated curves.

Figure 4 shows, for each model, the variation of the number of clusters, and the objective function, during the iterations of the proposed robust EM algorithm. We see that the number of clusters starting at $K = 100$ equal to the number of curves, rapidly decreases after the two first iterations and the majority of illegitimate clusters is discarded (more than 50%). Then, after few iteration, the algorithms converge to the actual number of clusters. We also see that in this case the penalized log-likelihood is increasing over the EM iterations and converges.

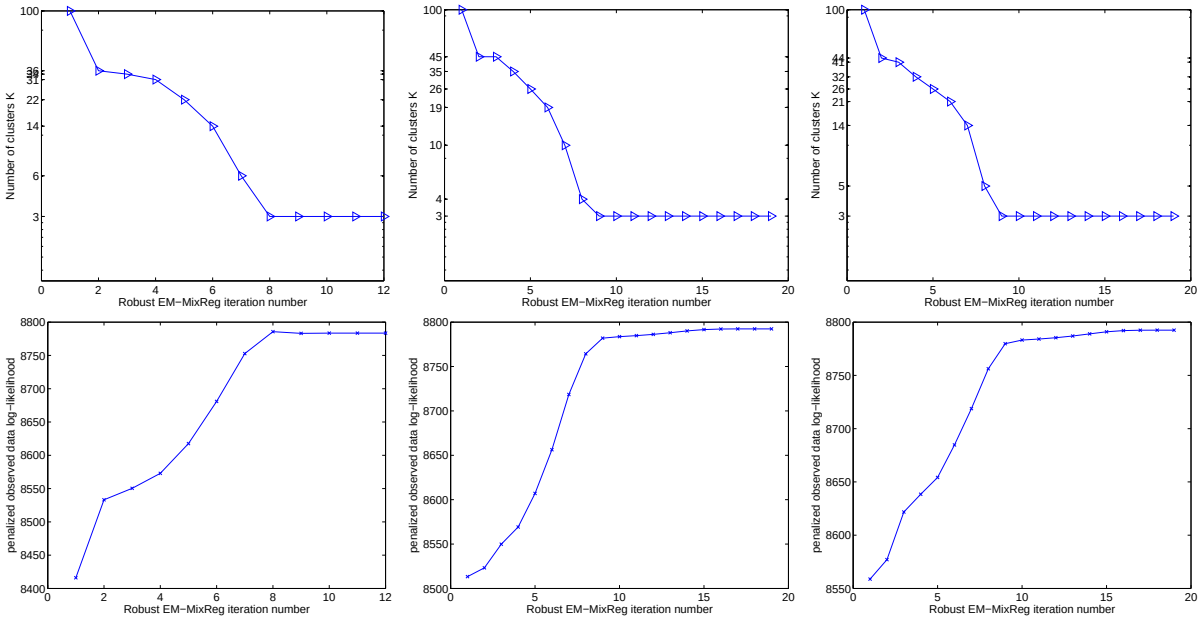
Number of clusters K
Robust EM−MixReg iteration number
penalized observed data log−likelihood
Robust EM−MixReg iteration number

Figure 4: Variation of the number of clusters and the value of the objective function during the iterations of the algorithm for the PRM (left), PSRM (middle) and PbSRM (right) for the simulated data.

To highlight the variation of the clustering over the iterations, in Figure 5, we show the clustering results obtained for the polynomial regression mixture model, at different iterations of proposed algorithm. We can see that after only one iteration, 64 illegitimate clusters are discarded. Then both the number of clusters and the clusters approximation are precisely updated until convergence. The convergence is achieved after only 12 iterations. It can be seen that for the second data set of non-linear curves, the proposed algorithm also provides accurate results. After starting with a number of clusters $K = 100$, the number of clusters decreases rapidly form 100 to 27 after only four iterations. Then the algorithm converges after 22 iterations and provides the actual number of clusters with precise clustering results.

4.1.2 Further simulation: Waveform curves of Brieman

In this experiment we consider the waveform data introduced in Breiman et al. (1984b) and studied in Hastie and Tibshirani (1996) and elsewhere. The waveform data consist in a three-class problem where each curve is generated as follows:

- $\mathbf{y}_1(t) = uh_1(t) + (1-u)h_2(t) + \epsilon_t$ for the class 1;

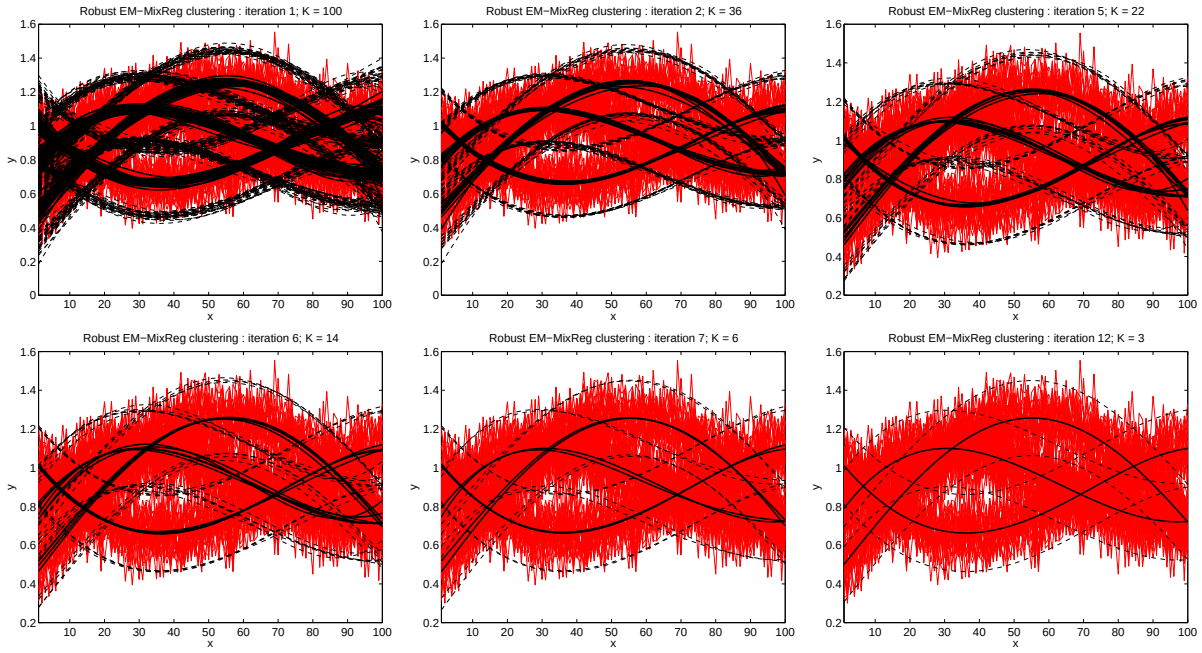
Robust EM−MixReg clustering : iteration 1; K = 100
Robust EM−MixReg clustering : iteration 2; K = 36
Robust EM−MixReg clustering : iteration 5; K = 22
Robust EM−MixReg clustering : iteration 6; K = 14
Robust EM−MixReg clustering : iteration 7; K = 6
Robust EM−MixReg clustering : iteration 12; K = 3
x
y

Figure 5: Clustering results for the simulated data in Fig. 1 obtained by the proposed algorithm for polynomial regressions mixtures at different iterations of the algorithm.

- $\mathbf{y}_2(t) = uh_2(t) + (1-u)h_3(t) + \epsilon_t$ for the class 2;
- $\mathbf{y}_3(t) = uh_1(t) + (1-u)h_3(t) + \epsilon_t$ for the class 3.

where u is a uniform random variable on $(0, 1)$,

- $h_1(t) = \max(6 - |t - 11|, 0)$;
- $h_2(t) = h_1(t - 4)$;
- $h_3(t) = h_1(t + 4)$.

and ϵ_t is a zero-mean Gaussian noise with unit standard deviation. The temporal interval considered for each curve is $[1; 21]$ with a constant period of sampling of 1 second. Figure 6 shows the waveform data mean functions from the generative model before the Gaussian noise is added and Fig. 7 shows a sample of 150 waveforms.

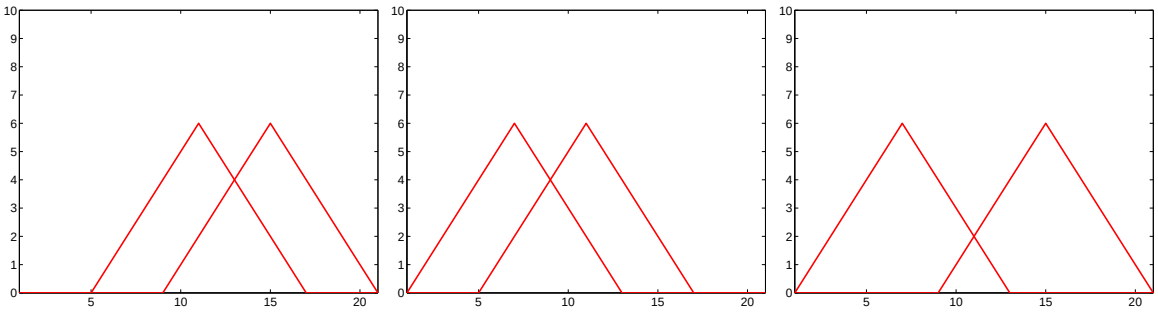

Figure 6: The waveform mean functions from the generative model before the Gaussian noise is added.

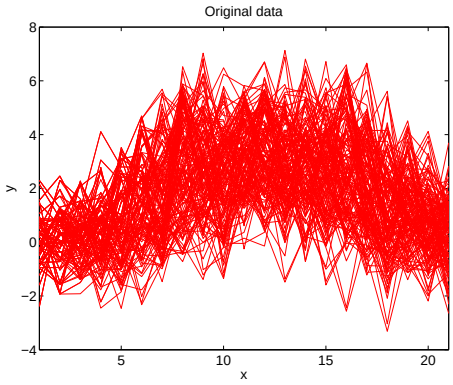

Figure 7: Waveform data.

Figures 8, 9 and 10 respectively show the obtained cluster results for the waveform data for the polynomial, spline and B-spline regression mixture. The number of clusters is correctly estimated by the three models. For this data, the spline regression models provide slightly better results in terms of clusters approximation than the polynomial regression mixture (here $p = 4$). This can be seen for the third cluster.

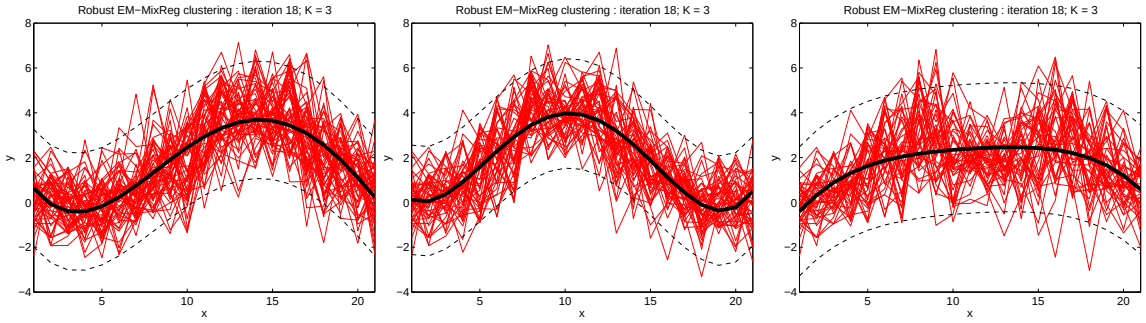
Robust EM−MixReg clustering : iteration 18; K = 3
Robust EM−MixReg clustering : iteration 18; K = 3
Robust EM−MixReg clustering : iteration 18; K = 3
x
y

Figure 8: Clustering results obtained by the proposed robust EM algorithm and the PRM (polynomial degree $p = 4$) model for the waveform data.

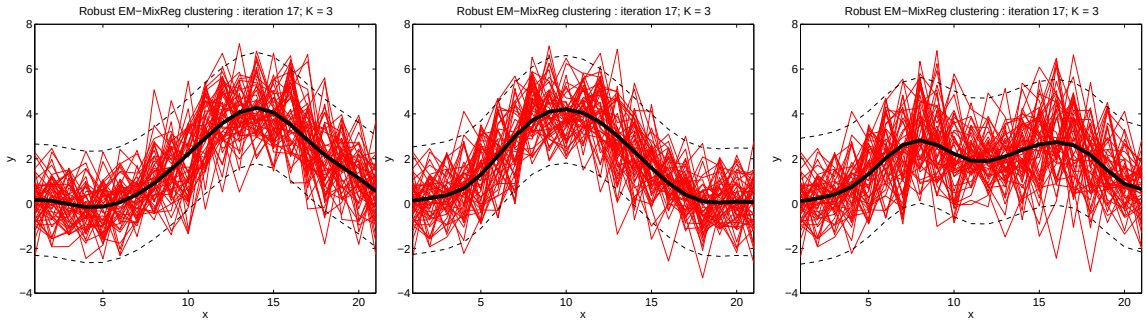
Robust EM−MixReg clustering : iteration 17; K = 3
Robust EM−MixReg clustering : iteration 17; K = 3
Robust EM−MixReg clustering : iteration 17; K = 3
x
y

Figure 9: Clustering results obtained by the proposed robust EM algorithm and the SRM with a cubic-spline of three knots for the waveform data.

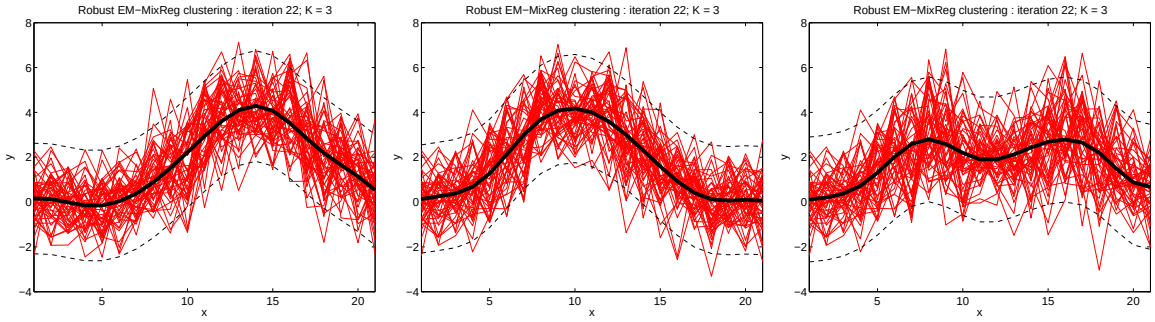
Robust EM−MixReg clustering : iteration 22; K = 3
Robust EM−MixReg clustering : iteration 22; K = 3
Robust EM−MixReg clustering : iteration 22; K = 3
x
y

Figure 10: Clustering results obtained by the proposed robust EM algorithm and the bSRM with a cubic b-spline of three knots for the waveform data.

Table 2 presents the clustering results averaged over 20 different sample of 500 curves. It includes the estimated number of clusters, the misclassification error rate and the absolute error between the true clusters proportions and variances and the estimated ones. One can observe that for all the models, the actual number of clusters is correctly retrieved. The misclassification error rate as well as the parameters estimation errors are slightly better for the spline regression models, in particular the B-spline regression mixture.

	actual	EM-PRM	EM-SRM	EM-bSRM
K	3	3	3	3
misc. error	-	$4.31 \pm (0.42)\%$	$2.94 \pm (0.88)\%$	$2.53 \pm (0.70)\%$
σ_1	1	$0.128 \pm (0.015)$	$0.108 \pm (0.015)$	$0.103 \pm (0.012)$
σ_2	1	$0.102 \pm (0.015)$	$0.090 \pm (0.011)$	$0.079 \pm (0.010)$
σ_3	1	$0.223 \pm (0.021)$	$0.180 \pm (0.014)$	$0.141 \pm (0.013)$
π_1	$\frac{1}{3}$	$0.0037 \pm (0.0018)$	$0.0035 \pm (0.0015)$	$0.0034 \pm (0.0015)$
π_2	$\frac{1}{3}$	$0.0029 \pm (0.0023)$	$0.0018 \pm (0.0015)$	$0.0012 \pm (0.0011)$
π_3	$\frac{1}{3}$	$0.0040 \pm (0.0062)$	$0.0037 \pm (0.0015)$	$0.0035 \pm (0.0014)$

Table 2: Clustering results for the waveform data.

In Figure 11, one can see the variation of the estimated number of clusters as well as the value of the objective function from one iteration to another for the three models. These results highlight the capability of the proposed algorithm to provide an accurate partition with a well adapted number of clusters.

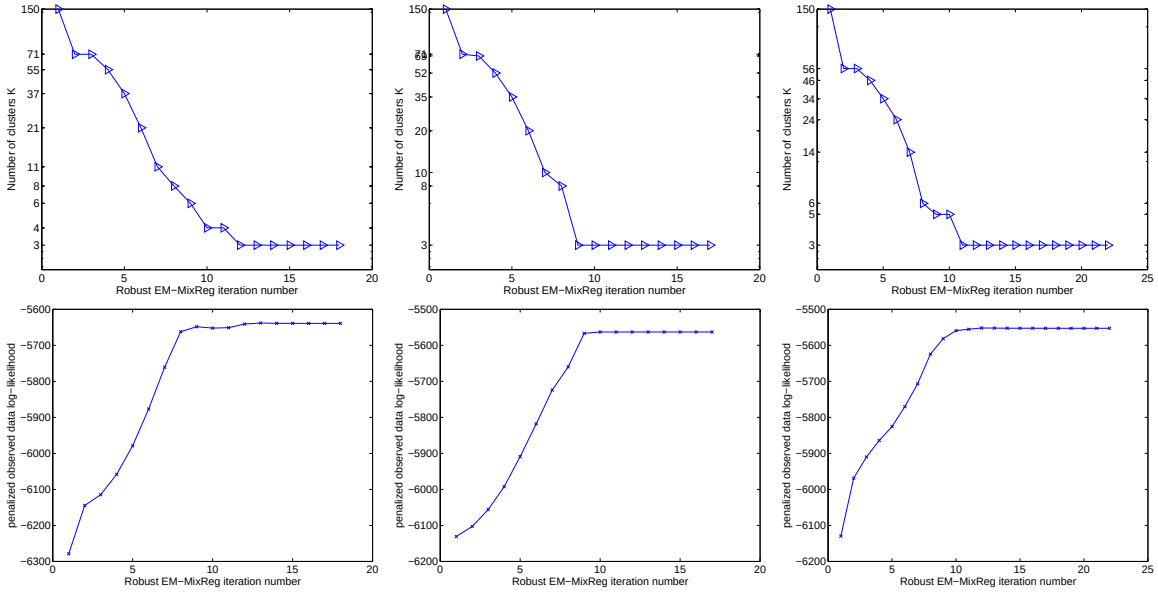
Number of clusters K
Robust EM−MixReg iteration number
penalized observed data log−likelihood
Robust EM−MixReg iteration number

Figure 11: Variation of the number of clusters and the value of the objective function during the iterations of the algorithm for the PRM (left), PSRM (middle) and PbSRM (right) for the waveform data.

4.2 Experiments on real data

In this section, we consider real data sets to evaluate the proposed approach. The considered data are curves issued from three different application domains: the phonemes data, the Yeast cell cycle data and the Topex/Poseidon satellite data.

4.2.1 Phonemes data

In this section, we use the phonemes dataset used in Ferraty and Vieu (2003)[1] which is a part of the original one available at `http://www-stat.stanford.edu/ElemStatLearn` and was described and used namely in (Hastie et al., 1995). The application context related to this dataset is a phoneme classification problem. The phonemes dataset correspond to log-periodograms y constructed from recordings available at different equispaced frequencies x for different phonemes. It contains five classes corresponding to the following five phonemes: "sh" as in "she", "dcl" as in "dark", "iy" as in "she", "aa" as in "dark" and "ao" as in "water". For each phoneme we have 400 log-periodograms at a 16-kHz sampling rate. We only retain the first 150 frequencies from each subject as in (Ferraty and Vieu, 2003). This dataset has been considered in a phoneme discrimination problem as in Hastie et al. (1995); Ferraty and Vieu (2003) where the aim is to

[1]Data from `http://www.math.univ-toulouse.fr/staph/npfda/`

predict the phoneme class for a new log-periodogram. Here we reformulate the problem into a clustering problem where the aim is to automatically group the phonemes data into classes. We assume that the cluster labels are missing. We also assume that the number of clusters is unknown. Thus, the proposed algorithm will be assessed in terms of estimating both the actual partition and the optimal number of clusters from the data. Fig. 12 shows the used 1000 log-periodograms (200 per cluster) and Fig. 13 shows the curves of the actual five phoneme classes, class by class.

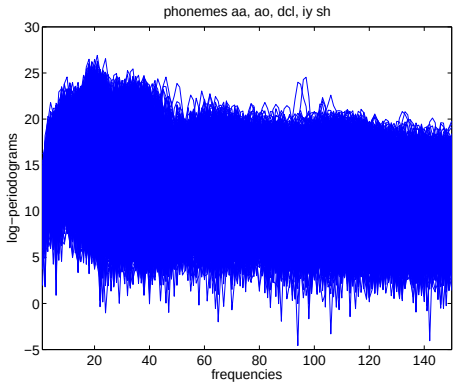

Figure 12: Phonemes data "ao", "aa", "yi", "dcl", "sh".

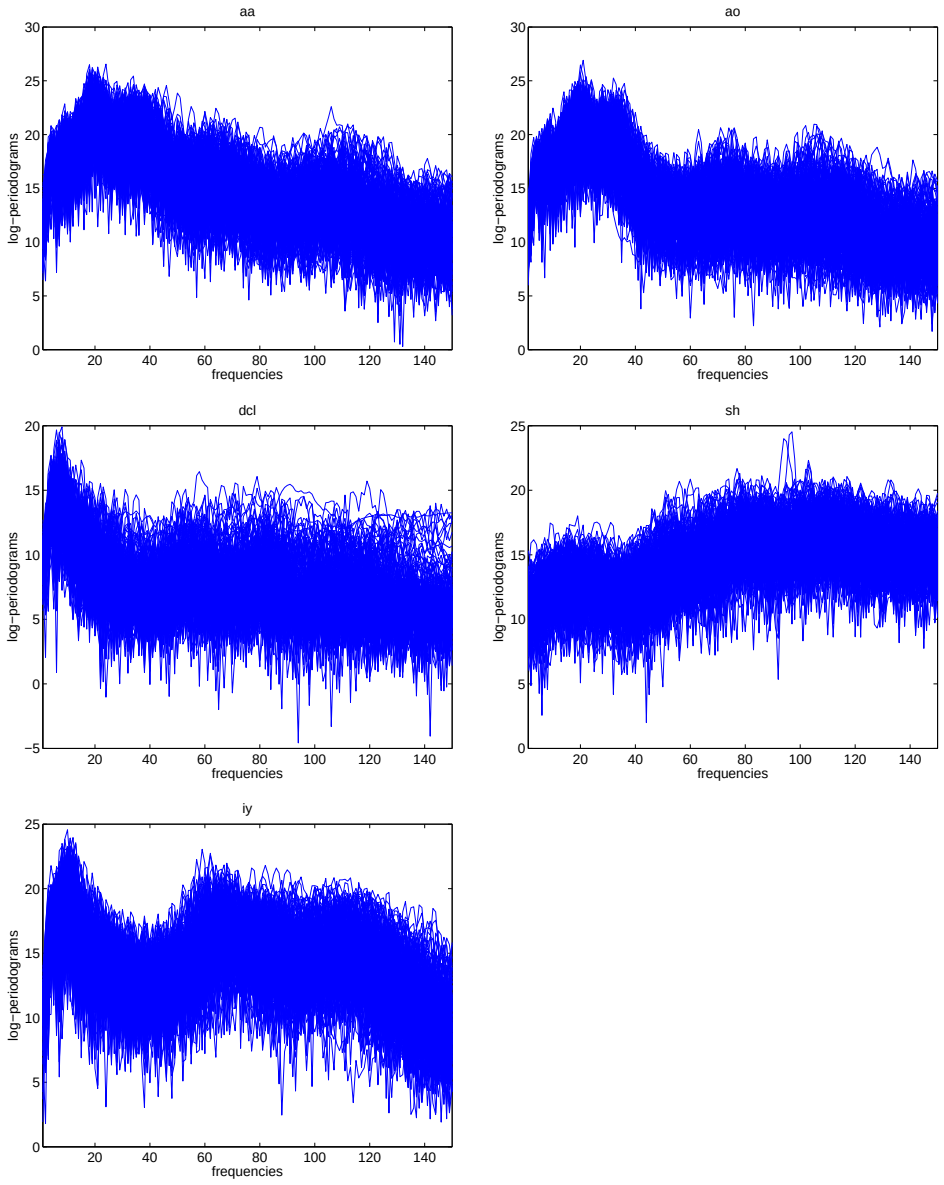

Figure 13: Phonemes data classes: "ao", "aa", "yi", "dcl", "sh".

Figures 14, 15 and 16 respectively show the obtained cluster results for the phonemes log-periodograms for the polynomial, spline and B-spline regression mixtures. The number of phoneme classes (five) is correctly estimated by the three models. We can also see that the spline regression models provide better results in terms of clusters approximation than the polynomial regression mixture (here $p = 7$). Notice that the value of $p = 7$ correspond to the polynomial regression mixture model with the best error rate for p varying from 4 to 8. The misclassification error rate in this case is 14.29 %.

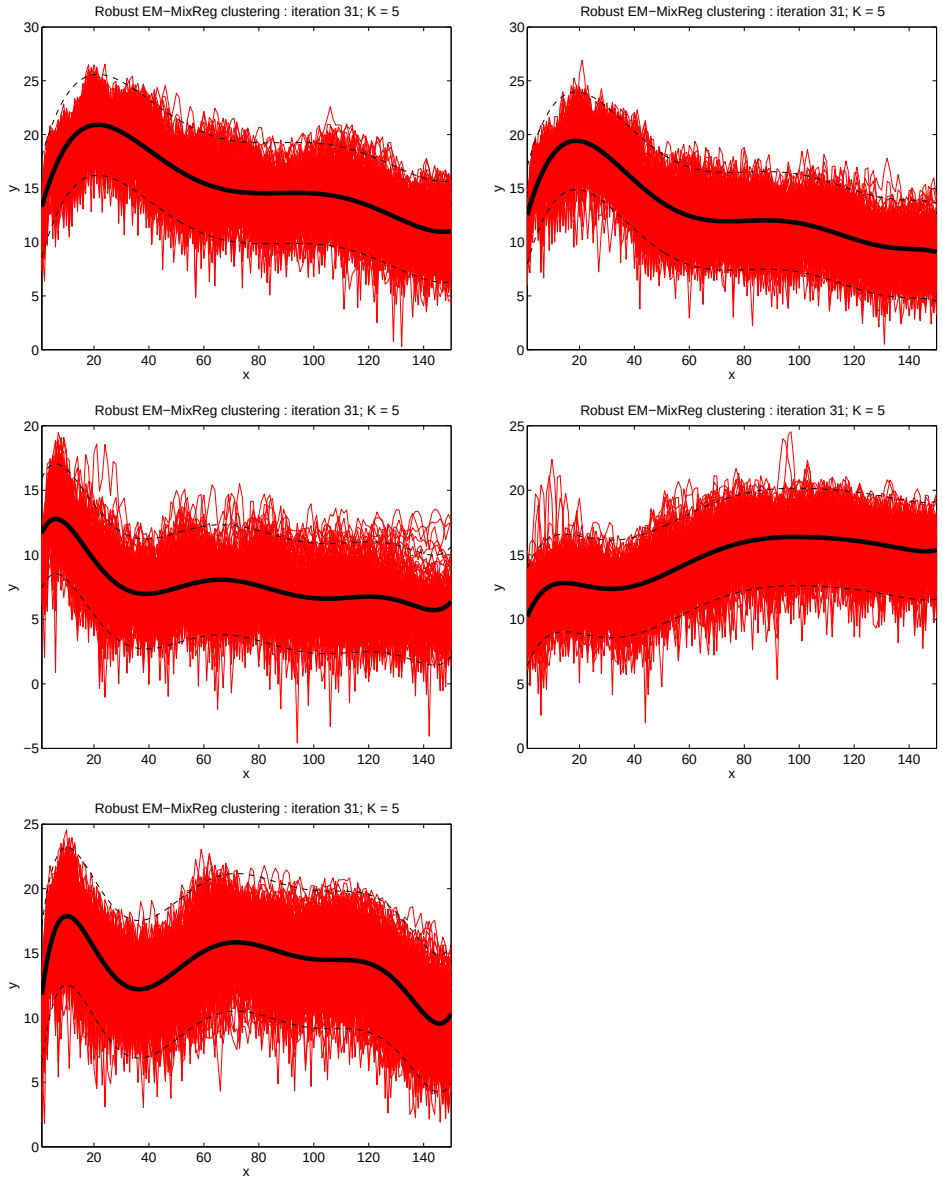
Robust EM−MixReg clustering : iteration 31; K = 5
Robust EM−MixReg clustering : iteration 31; K = 5
Robust EM−MixReg clustering : iteration 31; K = 5
Robust EM−MixReg clustering : iteration 31; K = 5
Robust EM−MixReg clustering : iteration 31; K = 5
x
y

Figure 14: Clustering results obtained by the proposed robust EM algorithm and the PRM model (polynomial degree $p = 7$) for the phonemes data.

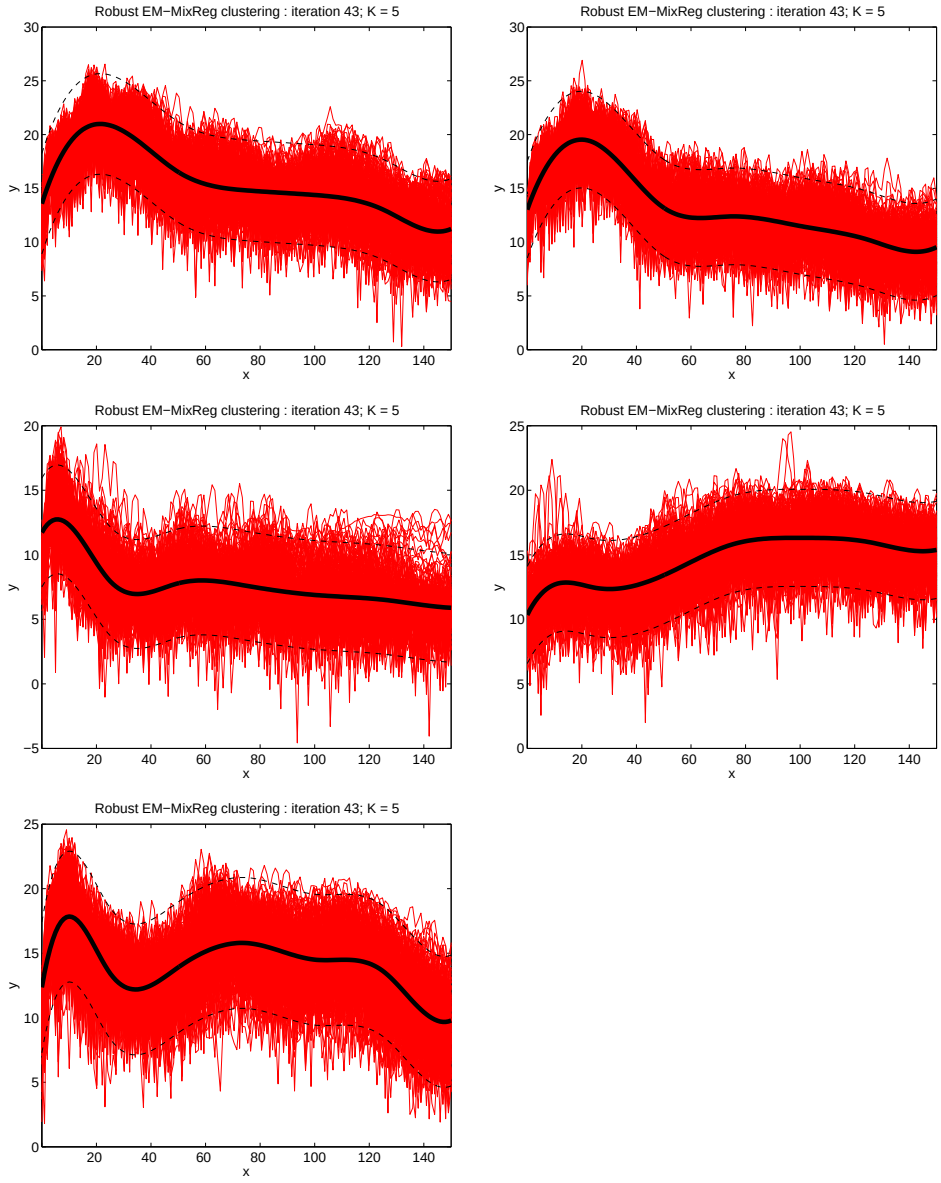
Robust EM−MixReg clustering : iteration 43; K = 5
Robust EM−MixReg clustering : iteration 43; K = 5
Robust EM−MixReg clustering : iteration 43; K = 5
Robust EM−MixReg clustering : iteration 43; K = 5
Robust EM−MixReg clustering : iteration 43; K = 5
x
y

Figure 15: Clustering results obtained by the proposed robust EM algorithm and the SRM model with a cubic B-spline of seven knots for the phonemes data.

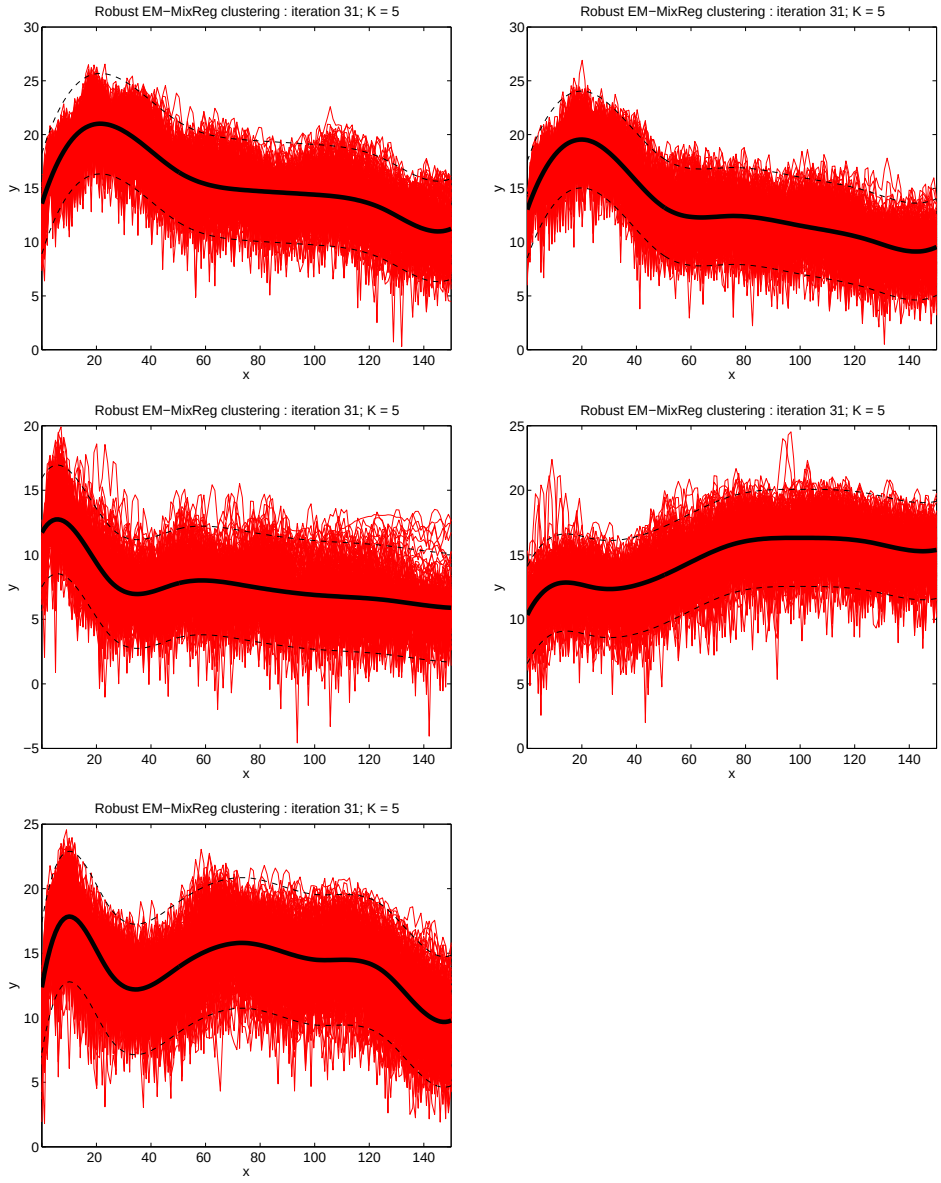
Robust EM−MixReg clustering : iteration 31; K = 5
Robust EM−MixReg clustering : iteration 31; K = 5
Robust EM−MixReg clustering : iteration 31; K = 5
Robust EM−MixReg clustering : iteration 31; K = 5
Robust EM−MixReg clustering : iteration 31; K = 5
x
y

Figure 16: Clustering results obtained by the proposed robust EM algorithm and the bSRM model with a cubic B-spline of seven knots for the phonemes data.

The values of the estimated number of clusters and the misclassification error rate corresponding to each of the three models are given in Table 3. One can see that the spline regressions mixture perform better than the polynomial regression mixture. In a general use of functional data modeling, the spline are indeed more adapted than simple polynomial modeling.

	EM-PRM	EM-SRM	EM-bSRM
$\hat{K}$	5	5	5
Misc. error rate	14.29 %	14.09 %	14.2 %

Table 3: Clustering results for the phonemes data.

In a similar way as previously, on Figure 17, one can see the variation of the estimated number of clusters as well as the value of the objective function as the learning proceeds. It can be observed that the number of clusters decreases very rapidly from 1000 to 51 for the polynomial regression mixture model, and to 44 for the spline and B-spline regression mixture models one iteration to another for the three models. The grand majority of illegitimate clusters is discarded at the beginning the learning process. Then, the number of clusters gradually decreases and the algorithm converge towards a partition with the actual number of clusters for the three models after at most 43 iterations. We can also see from the curve of the number of clusters and the objectives functions that the algorithm for the spline and B-spline regression mixture models behaves in a very similar way. We can also notice that the objective function becomes horizontal once the number of clusters is stabilized.

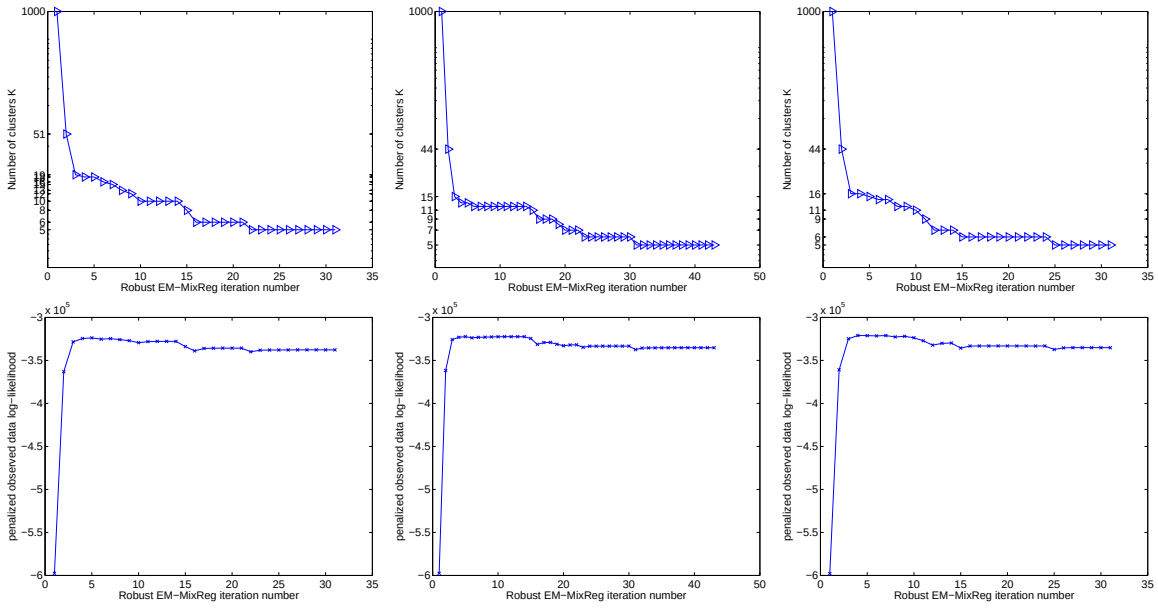

Figure 17: Variation of the number of clusters and the value of the objective function during the iterations of the algorithm for the PRM (left) PSRM (middle) and PbSRM (right) for the phonemes data.

4.2.2 Yeast cell cycle data

In this experiment, we consider the yeast cell cycle dataset (Cho et al., 1998). The original yeast cell cycle data represent the fluctuation of expression levels of approximately 6000 genes over 17 time points corresponding to two cell cycles (Cho et al., 1998). This data set has been used to demonstrate effectiveness of clustering techniques for for time course Gene expression data in

bio-informatics such as model-based clustering as inYeung et al. (2001). We used the standardized subset constructed by Yeung et al. (2001) available in `http://faculty.washington.edu/kayee/model/`[2]. This dataset referred to as the subset of the 5-phase criterion in Yeung et al. (2001) contains 384 genes expression levels over 17 time points. The utility of the cluster analysis in this case is therefore to reconstruct this five class partition. Fig. 18 shows the 384 curves of the yeast cell cycle data and Fig. 19 shows the curves of each of five clusters.

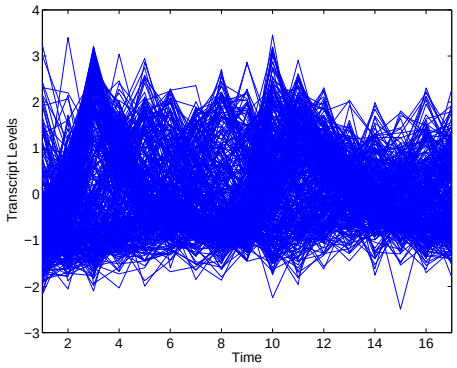

Figure 18: Yeast cell cycle data.

[2]The complete data set was downloaded from `http://genome-www.stanford.edu/cellcycle/`.

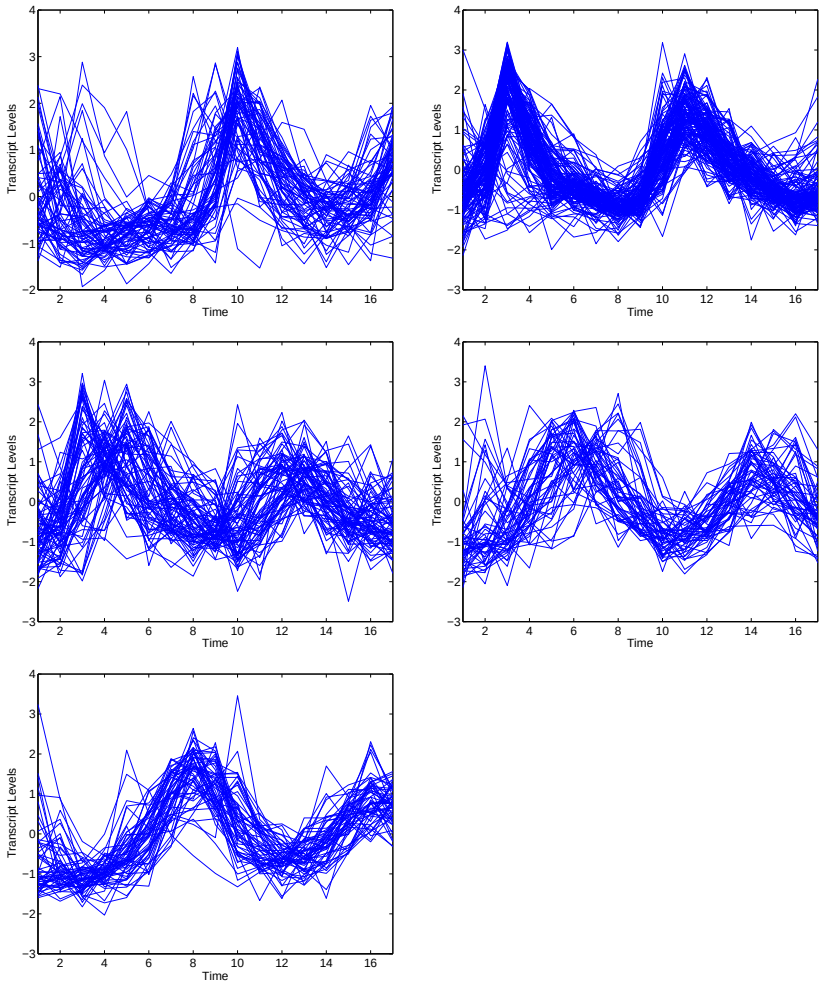

Figure 19: The five actual clusters of the used yeast cell cycle data.

The clustering results are shown in Figures 20, 21 and 22 respectively for the polynomial, spline and B-spline regression mixture models. Both the PRM model and the SRM provide similar partitions with four clusters. The actual clusters in middle left and bottom in Figure 19 look to be merged into the cluster in middle left in Figure 20 (respectively in Figure 21). However, the bSRM model infers the actual number of clusters. The Rand index for the obtained partition is 0.7914.

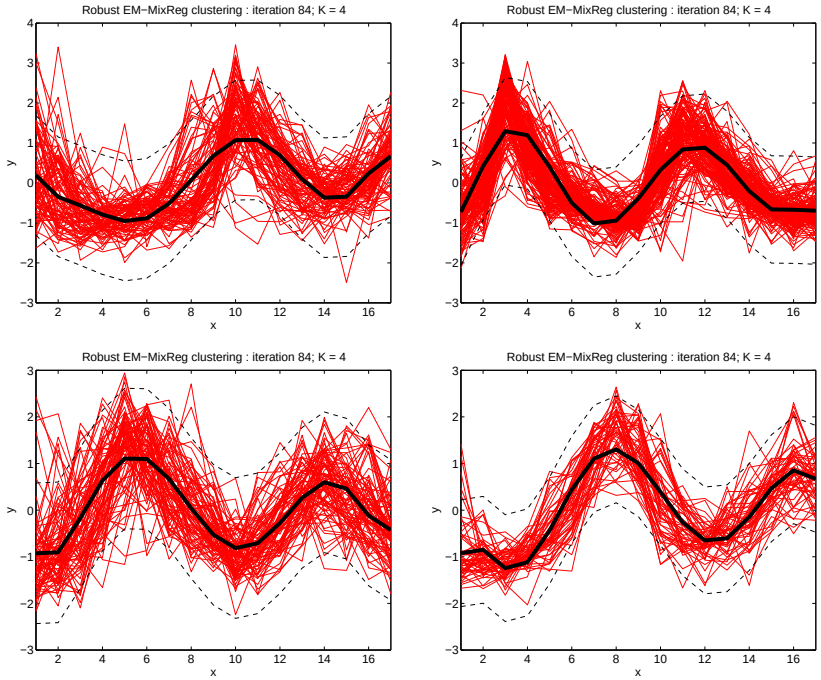
Robust EM−MixReg clustering : iteration 84; K = 4
Robust EM−MixReg clustering : iteration 84; K = 4
Robust EM−MixReg clustering : iteration 84; K = 4
Robust EM−MixReg clustering : iteration 84; K = 4
x
y

Figure 20: Clustering results obtained by the proposed robust EM algorithm and the PRM model (polynomial degree $p = 6$) for the yeast cell cycle data.

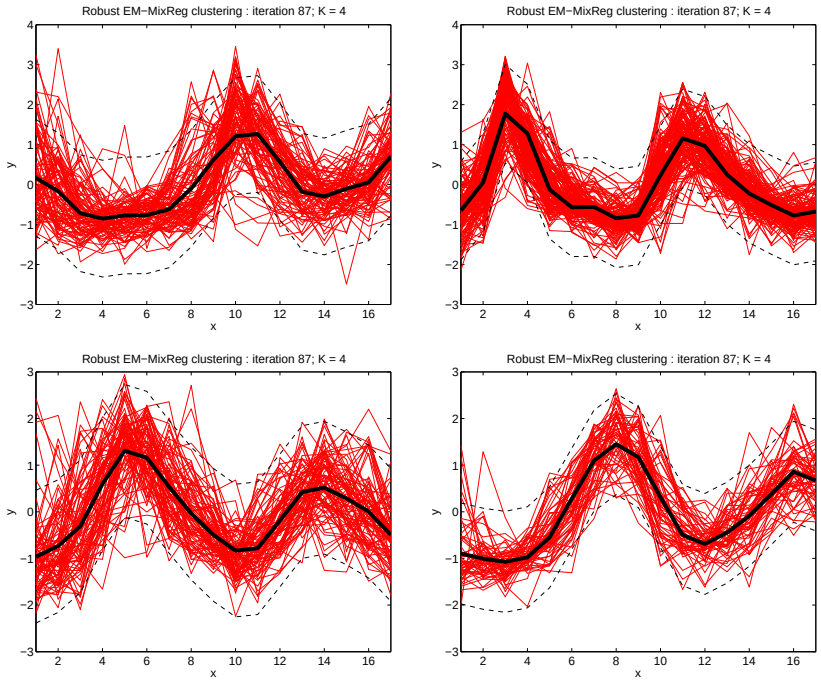
Robust EM−MixReg clustering : iteration 87; K = 4
Robust EM−MixReg clustering : iteration 87; K = 4
Robust EM−MixReg clustering : iteration 87; K = 4
Robust EM−MixReg clustering : iteration 87; K = 4
x
y

Figure 21: Clustering results obtained by the proposed robust EM algorithm and the SRM model with a cubic spline of 7 knots for the yeast cell cycle data.

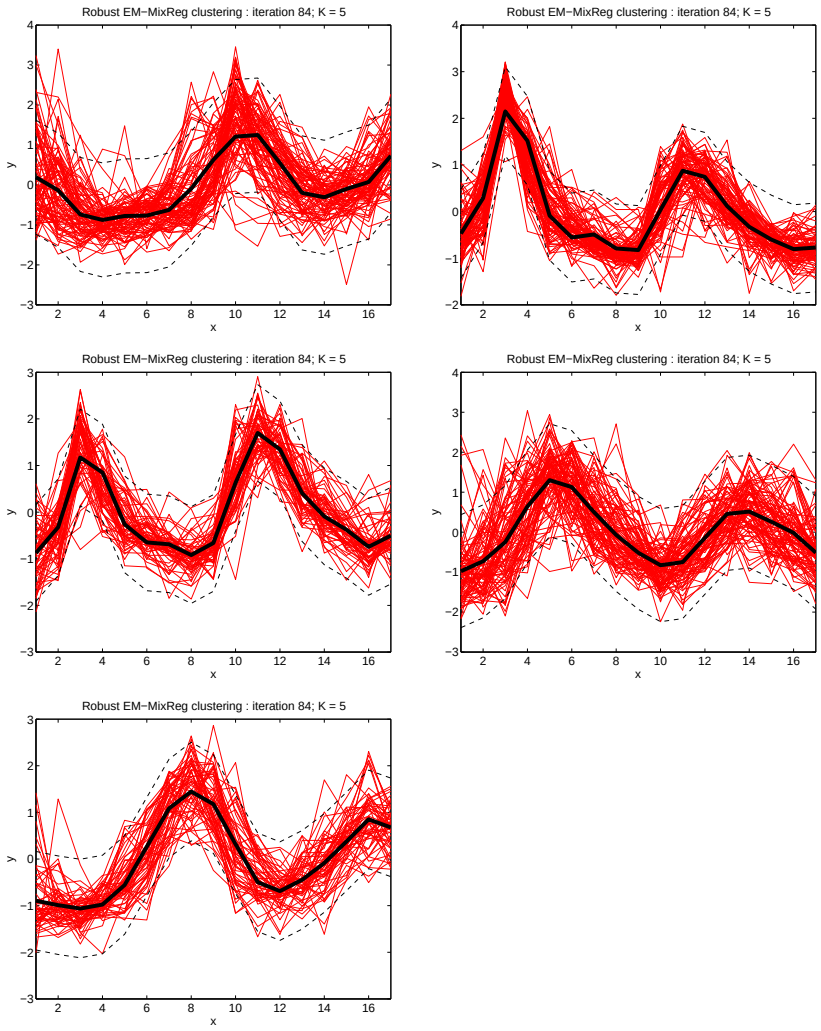
Robust EM−MixReg clustering : iteration 84; K = 5
Robust EM−MixReg clustering : iteration 84; K = 5
Robust EM−MixReg clustering : iteration 84; K = 5
Robust EM−MixReg clustering : iteration 84; K = 5
Robust EM−MixReg clustering : iteration 84; K = 5
x
y

Figure 22: Clustering results obtained by the proposed robust EM algorithm and the bSRM model with a cubic B-spline of 7 knots for the yeast cell cycle data.

Figure 23 shows the variation of the number of clusters and the value of the objective function during the iterations of the algorithm for three models. We can see that the number of clusters starts with $n = 384$ clusters and more than half is discarded after one iteration. Then it gradually decreases and stabilized until convergence. The shape of the objective function also becomes horizontal when it is converged.

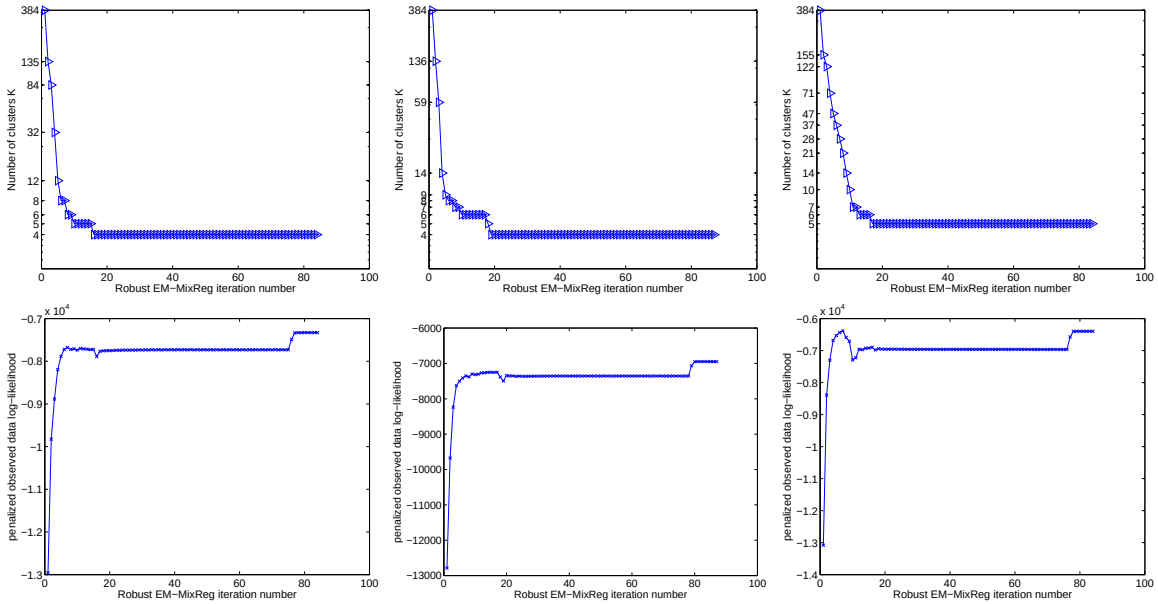

Figure 23: Variation of the number of clusters and the value of the objective function during the iterations of the algorithm for the PRM (left) PSRM (middle) and PbSRM (right) for the yeast cell cycle data.

4.2.3 Topex/Poseidon satellite data

The last used real dataset is the Topex/Poseidon radar satellite data namely used in Dabo-Niang et al. (2007); Hébrail et al. (2010) [3] This dataset were registered by the satellite Topex/Poseidon around an area of 25 kilometers upon the Amazon River. The data contain $n = 472$ waveforms of the measured echoes, sampled at $m = 70$ (number of echoes). The actual number of clusters and the actual partition are unknown for this dataset. The curves of this data set are shown in Figure 24.

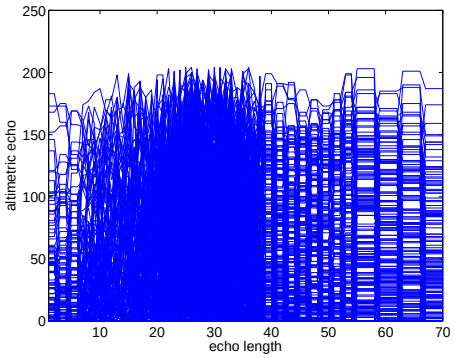

Figure 24: Topex/Poseidon satellite curves.

Figures 25, 26 and 27 show the obtained clustering results. First, one can see that the provided

[3]Satellite data are available at `http://www.lsp.ups-tlse.fr/staph/npfda/npfda-datasets.html`.

solution for the polynomial regression mixture fitting is rather an overall rough approximation that provides three clusters. Even they don't look similar, the polynomial fitting for this type of curves is not adapted. This is because the curves namely present peaks and transitions. The solutions provided by the spline and B-spline regression mixture is more informative about the underlying structure of this dataset. We indeed used a linear (B)-spline for this dataset in order to allow piecewise linear function approximation and thus to better recover the possible peaks and transitions in the curves. As a result, both the SRM and the bSRM provide a five class partition. The partitions are quasi-identical and contain clearly informative clusters. We can indeed see different forms of waves that summarize the general structure underlying of this data. We can see two clusters containing curves presenting one narrow peak (the cluster in top left and the one in bottom). The two clusters differ with the peak location in x. The cluster in middle left contains curves with one less narrow peak. The first cluster contains curves that look to have two large peaks. Finally, the third cluster in middle left looks to contain curves without peaks and with a part rather flat. Furthermore, we can see that the structure more clear with the cluster mean (prototypes) than with the raw curves. The spline regression mixture models thus helps to better understand the underlying structure of the data as well as to recover a plausible number of clusters from the data. In addition, the found number of cluster (five) also equals the one found by Dabo-Niang et al. (2007) by using another hierarchical nonparametric kernel-based unsupervised classification technique. The mean curves for the five terminal groups reflecting the hidden structure provided by the proposed approach for both the PRM and the bSRM are similar to those in Dabo-Niang et al. (2007). On the other hand, one can also see that this result is similar to the one found in Hébrail et al. (2010), most of the profiles are present in the two results. The slight difference can be attributed to the fact that the result in Hébrail et al. (2010) is provided from a two-stage scheme which includes and additional pre-clustering step using the Self Organizing Map (SOM), rather by directly applying the piecewise regression model to the raw data. We also notice that, in Hébrail et al. (2010), the number of clusters was set to twenty and the clustering procedure was two-fold. The authors indeed used a two-fold scheme and first performed a topographic clustering step using the SOM, and then apply a K-means-like approach to the results of the SOM. However, in our approach, we directly apply the proposed algorithm to the raw satellite data without a preprocessing step. In addition, the number of clusters is automatically inferred from the data. We also can observe that, the found five clusters here do summarize the general behavior of the twenty clusters in Hébrail et al. (2010) which can be summarized in clusters with one narrow shifted peak, less narrow peak, two large peaks and finally a cluster containing curves with sharp increase followed by a slow decrease.

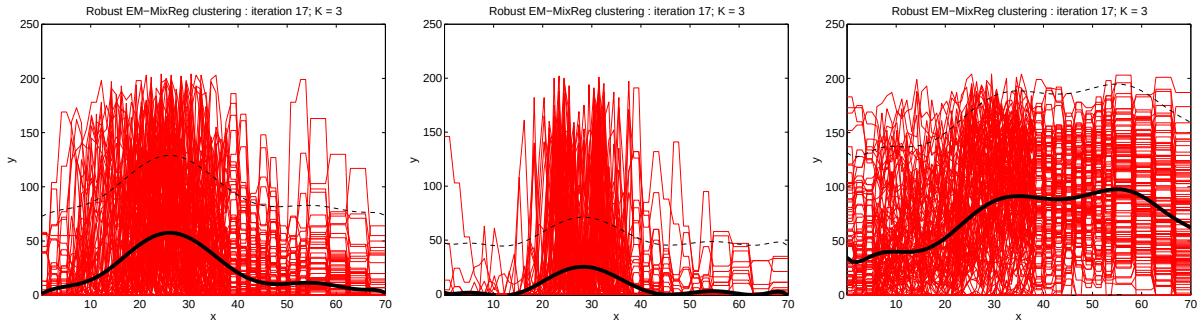
Robust EM−MixReg clustering : iteration 17; K = 3
Robust EM−MixReg clustering : iteration 17; K = 3
Robust EM−MixReg clustering : iteration 17; K = 3
y
x

Figure 25: Clustering results obtained by the proposed robust EM algorithm and the PRM with $p = 9$ for the satellite data.

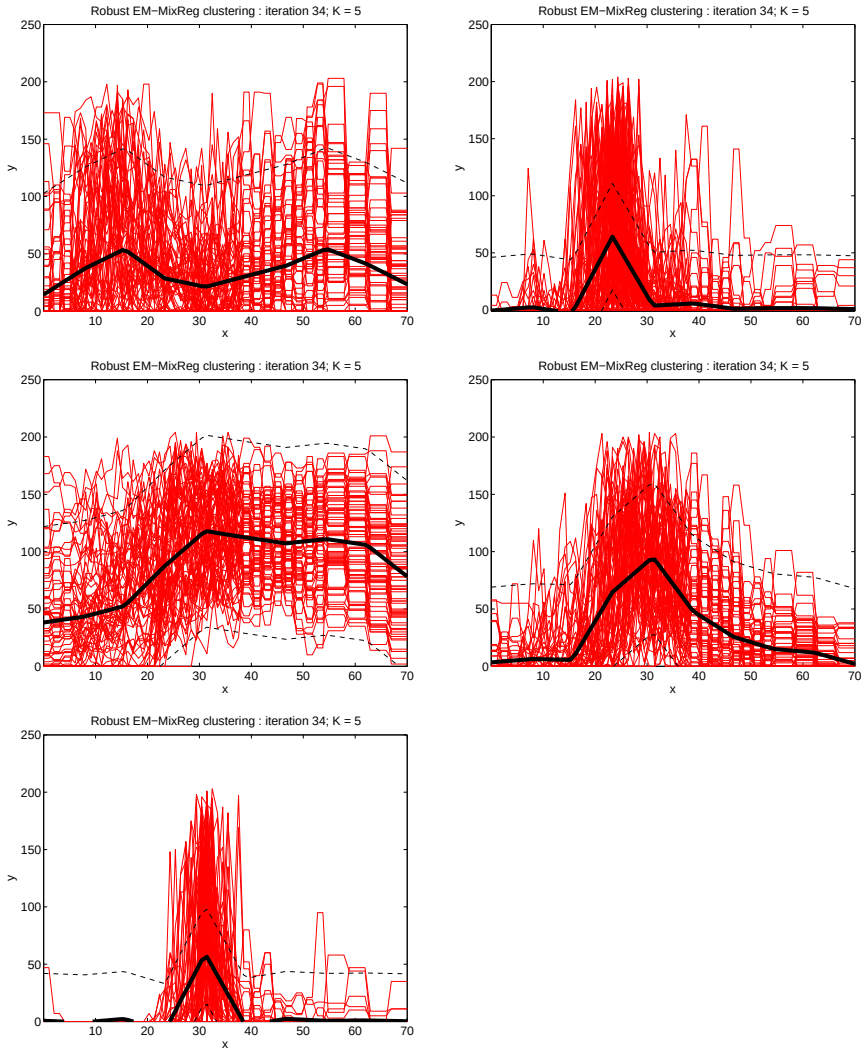
Robust EM−MixReg clustering : iteration 34; K = 5
Robust EM−MixReg clustering : iteration 34; K = 5
Robust EM−MixReg clustering : iteration 34; K = 5
Robust EM−MixReg clustering : iteration 34; K = 5
Robust EM−MixReg clustering : iteration 34; K = 5
y
x

Figure 26: Clustering results obtained by the proposed robust EM algorithm and the SRM model with a linear spline of 8 knots for the satellite data.

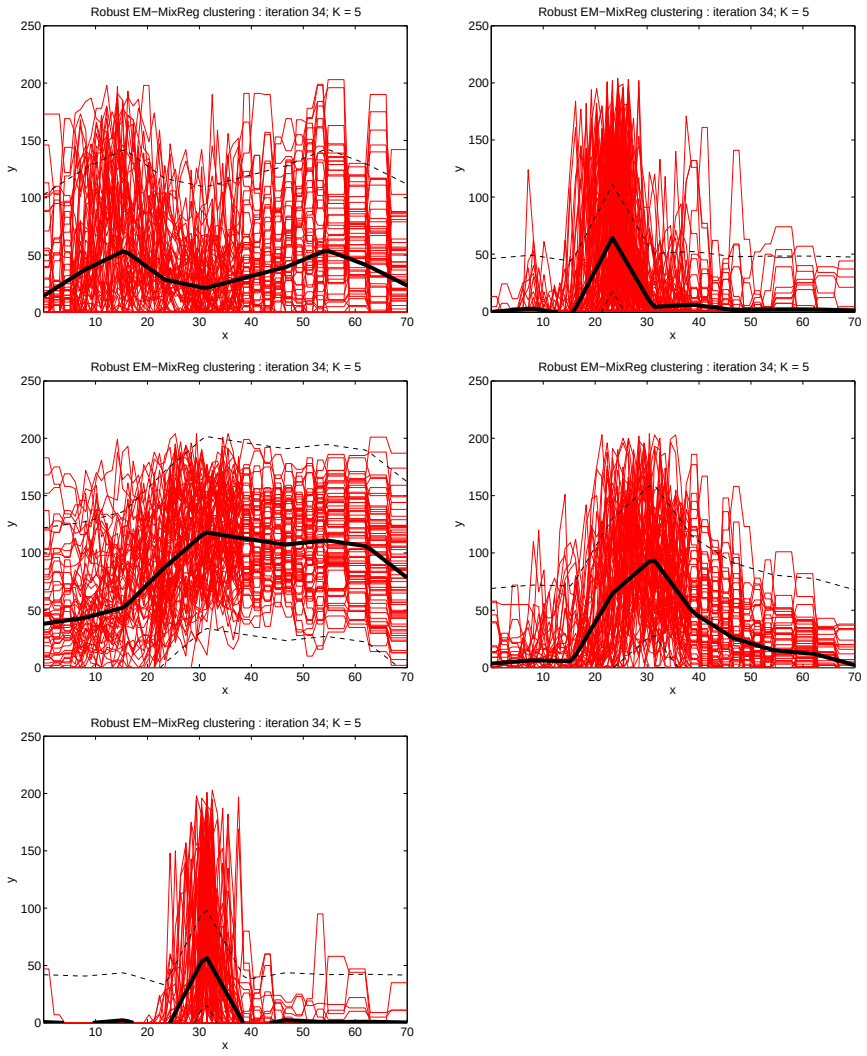
Robust EM−MixReg clustering : iteration 34; K = 5
Robust EM−MixReg clustering : iteration 34; K = 5
Robust EM−MixReg clustering : iteration 34; K = 5
Robust EM−MixReg clustering : iteration 34; K = 5
Robust EM−MixReg clustering : iteration 34; K = 5
x
y

Figure 27: Clustering results obtained by the proposed robust EM algorithm and the bSRM model with a linear B-spline of 8 knots for the satellite data.

Figure 28 shows that the algorithms converge after at most 35 iterations. The variation of the number of clusters during the iterations of the algorithm shows that after starting with 472 (n) clusters, the number of clusters rapidly decreases to 59 for the PRM and to 95 for both the SRM and the bSRM models. Then it gradually decreases until the number of clusters is stabilized. The variation of the value of the objective function during the iterations of the algorithm also shows that it becomes horizontal at converge which correspond to the stabilization of the partition.

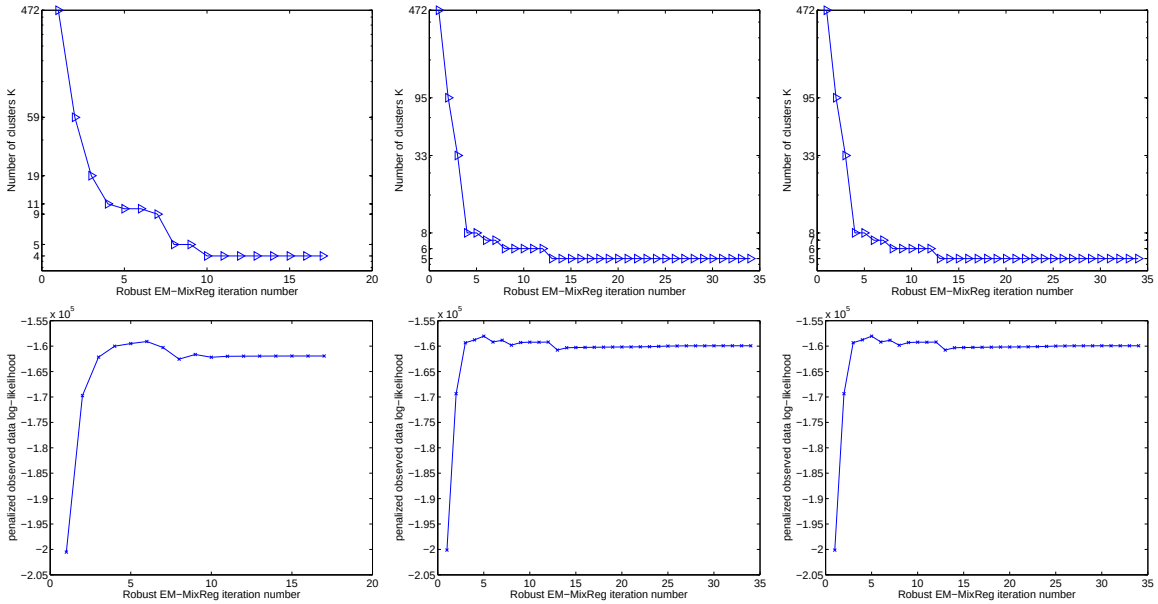

Figure 28: Variation of the number of clusters and the value of the objective function during the iterations of the algorithm for the PRM (left) PSRM (middle) and PbSRM (right) for the satellite data.

5 Conclusions and discussion

In this paper, we presented a new robust EM algorithm for model-based clustering using regression mixtures. It optimizes a penalized observed-data log-likelihood and overcomes both the problem of sensitivity to initialization and determining the optimal number of clusters for standard EM for regression mixtures. The experimental results on simulated data and real-world data demonstrate the benefit of the proposed approach for applications in curve clustering.

We note that, for the polynomial regression mixture, the choice of the polynomial degree can be performed in such a way to obtain the best partition. In practice, we varied the polynomial degree from 3 to 7 for the simulated data and from 4 to 7 for the waveform data. The obtained clustering results was closely similar and the number of clusters was always correctly selected. For the phonemes data and the yeast cell cycle data, the polynomial degree with the best solution was retained. However, for a more general use in functional data clustering and approximation, the splines are clearly more adapted. In practice, for the spline and B-spline regression mixtures, we used cubic (B)-splines because cubic splines, which correspond to a spline of order 4 which are are sufficient to approximate smooth functions. However, when the data present irregularity, such as a kind of piecewise non continuous functions, which is the case of the the Topex/poseidon sateliite data, we use a linear (B)-spline approximation. We also note that the algorithm is fast for the three models and converges after a few number of iterations, and took at most less than one 45 seconds for the phonemes data. For the other data, it took only few seconds. This makes is useful for real practical situations.

In this paper, we considered the problem of unsupervised fitting of regression mixtures with

unknown number of components. The regression mixture models are similar to the mixture of experts (ME) model Jacobs et al. (1991). Although similar, mixture of experts differ from curve clustering models in many respects. One of the main differences is that the ME model consists in a fully conditional mixture while in the regression mixture, only the component densities are conditional. Indeed, the mixing proportions are constant for the regression mixture, while in the mixture of experts, they mixing proportions (known as the gating functions of the network) are modeled as a function of the inputs, generally as a logistic or a softmax function. One interesting future direction is to extend the proposed approach to the problem of fitting mixture of experts (Jacobs et al., 1991) and hierarchical mixture of experts (Jordan and Jacobs, 1994) with unknown number of experts.

A Construction of B-splines basis functions

Given the sequence of knots $\xi_0 < \xi_1 < \ldots < \xi_{L+1}$ (ξ_0 and ξ_{L+1} are the two bounds of x), let us define the augmented knot sequence ζ such that

- $\zeta_1 \leq \zeta_2 \ldots \leq \zeta_M \leq \xi_0$;
- $\zeta_{M+\ell} = \xi_\ell,\ \ell = 1, \ldots, L$;
- $\xi_{L+1} \leq \zeta_{L+M+1} \leq \zeta_{KL+M+2} \ldots \leq \zeta_{L+2M}$.

The actual values of these additional knots beyond the boundary are arbitrary, and a common choice is to make them all the same and equal to ξ_0 and ξ_{L+1} respectively. Let us denote by $B_{\ell,M}(t)$ the ℓth B-spline basis function of order M for the knot-sequence

$$\zeta_1 \leq \zeta_2 \ldots \leq \zeta_M \leq \xi_0 < \xi_1 < \ldots < \xi_L < \xi_{L+1} \leq \zeta_{L+M+1} \leq \zeta_{L+M+2} \ldots \leq \zeta_{L+2M}.$$

These basis functions are defined recursively as follows:

- $B_{\ell,1}(x_{ij}) = \mathbb{1}_{[\zeta_j, \zeta_{j+1}]},\ \forall \ell = 1, \ldots, L + 2M - 1$;
- $B_{\ell,M}(x_{ij}) = \frac{x_{ij} - \zeta_\ell}{\xi_{\ell+M-1} - \zeta_\ell} B_{\ell,M-1}(x_{ij}) + \frac{\zeta_{\ell+M} - x_{ij}}{\zeta_{\ell+M} - \zeta_{\ell+1}} B_{\ell+1,M-1}(x_{ij}),\ \forall \ell = 1, \ldots, L + M.$

For the B-spline regression model, the jth row $\mathbf{b}_j$ ($j = 1, \ldots, m_i$) of the $m_i \times (L + M)$ B-spline regression matrix $\mathbf{B}_i$ for the ith curve is then constructed as follows:

$$\mathbf{b}_j = [B_{1,M}(x_{ij}),\ B_{2,M}(x_{ij}), \ldots,\ B_{L+M,M}(x_{ij})].$$

B Estimation of the mixing proportion

Consider the problem of finding the maximum of the function (27)

$$Q_\pi(\lambda, \pi_1, \ldots, \pi_K; \boldsymbol{\theta}^{(q)}) = \sum_{i=1}^{n} \sum_{k=1}^{K} \tau_{ik}^{(q)} \log \pi_k + \lambda n \sum_{k=1}^{K} \pi_k \log \pi_k \tag{33}$$

w.r.t the mixing proportions $(\pi_1, \ldots, \pi_K)$ subject to the constraint $\sum_{k=1}^{K} \pi_k = 1$. To perform this constrained maximization, we introduce the Lagrange multiplier α such that the resulting Lagrangian function is given by:

$$L(\pi_1, \ldots, \pi_K) = \sum_{i=1}^{n} \sum_{k=1}^{K} \tau_{ik}^{(q)} \log \pi_k + \lambda n \sum_{k=1}^{K} \pi_k \log \pi_k + \alpha(1 - \sum_{k=1}^{K} \pi_k). \tag{34}$$

Taking the derivatives of the Lagrangian with respect to π_k for $k = 1, \ldots, K$ we obtain:

$$\frac{\partial L(\pi_1, \ldots, \pi_K)}{\partial \pi_k} = \frac{\sum_{i=1}^{n} \tau_{ik}^{(q)}}{\pi_k} + \left(\lambda \sum_{i=1}^{n} (\log \pi_k + 1) \right) - \alpha \tag{35}$$

Then, setting these derivatives to zero yields:

$$\frac{\sum_{i=1}^{n} \tau_{ik}^{(q)}}{\pi_k} + n\lambda \log \pi_k + n\lambda = \alpha \tag{36}$$

By multiplying each hand side of (36) by π_k and summing over k we get

$$\sum_{k=1}^{K} \pi_k \times \left(\frac{\sum_{i=1}^{n} \tau_{ik}^{(q)}}{\pi_k} + n\lambda \log \pi_k + n\lambda \right) = \sum_{k=1}^{K} \alpha \times \pi_k \tag{37}$$

which implies that

$$n + n\lambda \sum_{k=1}^{K} \pi_k \log \pi_k + n\lambda = \alpha\cdot \tag{38}$$

Then, from (36) we can write

$$\sum_{i=1}^{n} \tau_{ik}^{(q)} + n\lambda\pi_k \log \pi_k + n\lambda\pi_k = n\pi_k + n\lambda\pi_k \sum_{h=1}^{K} \pi_h \log \pi_h + n\lambda\pi_k \tag{39}$$

$$n\pi_k = \sum_{i=1}^{n} \tau_{ik}^{(q)} + n\lambda\pi_k \log \pi_k - n\lambda \sum_{h=1}^{K} \pi_h \log \pi_k \tag{40}$$

Finally we therefore get the updating formula for the mixing proportions π_k's, that is

$$\pi_k^{(q+1)} = \frac{\sum_{i=1}^{n} \tau_{ik}^{(q)}}{n} + \lambda \pi_k^{(q)} \left(\log \pi_k^{(q)} - \sum_{h=1}^{K} \pi_h^{(q)} \log \pi_h^{(q)} \right) \tag{41}$$

$\forall k \in \{1, \ldots, K\}$.

References

H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974.

J. D. Banfield and A. E. Raftery. Model-based gaussian and non-gaussian clustering. *Biometrics*, 49(3):803–821, 1993.

C. Biernacki, G. Celeux, and G Govaert. Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(7):719–725, 2000.

C. Biernacki, G. Celeux, and G. Govaert. Choosing starting values for the EM algorithm for getting the highest likelihood in multivariate gaussian mixture models. *Computational Statistics and Data Analysis*, 41:561–575, 2003.

L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification And Regression Trees*. Wadsworth, New York, 1984a.

L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification And Regression Trees*. Wadsworth, New York, 1984b.

G. Celeux and J. Diebolt. The SEM algorithm a probabilistic teacher algorithm derived from the EM algorithm for the mixture problem. *Computational Statistics Quarterly*, 2(1):73–82, 1985.

G. Celeux and G. Govaert. A classification EM algorithm for clustering and two stochastic versions. *Computational Statistics and Data Analysis*, 14:315–332, 1992.

F. Chamroukhi. *Hidden process regression for curve modeling, classification and tracking*. Ph.D. thesis, Université de Technologie de Compiègne, Compiègne, France, 2010.

F. Chamroukhi, A. Samé, G. Govaert, and P. Aknin. Time series modeling by a regression approach based on a latent process. *Neural Networks*, 22(5-6):593–602, 2009.

F. Chamroukhi, A. Samé, G. Govaert, and P. Aknin. A hidden process regression model for functional data description. application to curve discrimination. *Neurocomputing*, 73(7-9): 1210–1221, March 2010.

F. Chamroukhi, A. Samé, P. Aknin, and G. Govaert. Model-based clustering with hidden markov model regression for time series with regime changes. In *Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN)*, pages 2814–2821, 2011.

Raymond J. Cho, Michael J. Campbell, Elizabeth A. Winzeler, Lars Steinmetz, Andrew Conway, Lisa Wodicka, Tyra G. Wolfsberg, Andrei E. Gabrielian, David Landsman, David J. Lockhart, and Ronald W. Davis. A genome-wide transcriptional analysis of the mitotic cell cycle. *Molecular Cell*, 2(1):65–73, 1998.

Sophie Dabo-Niang, Frédéric Ferraty, and Philippe Vieu. On the using of modal curves for radar waveforms classification. *Computational Statistics & Data Analysis*, 51(10):4878 – 4890, 2007.

C. Deboor. *A Practical Guide to Splines*. Springer-Verlag, 1978.

A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of The Royal Statistical Society, B*, 39(1):1–38, 1977.

Wayne DeSarbo and William Cron. A maximum likelihood methodology for clusterwise linear regression. *Journal of Classification*, 5(2):249–282, 1988.

Susana Faria and Gilda Soromenho. Fitting mixtures of linear regressions. *Journal of Statistical Computation and Simulation*, 80(2):201–225, February 2010.

Frédéric Ferraty and Philippe Vieu. Curves discrimination: a nonparametric functional approach. *Computational Statistics & Data Analysis*, 44(1-2):161–173, 2003.

Mário A. T. Figueiredo and Anil K. Jain. Unsupervised learning of finite mixture models. *IEEE Transactions Pattern Analysis and Machine Intelligence*, 24:381–396, 2000.

C. Fraley and A. E. Raftery. How many clusters? which clustering method? answers via model-based cluster analysis. *The Computer Journal*, 41(8):578–588, August 1998.

C. Fraley and A. E. Raftery. Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association*, 97:611–631, 2002.

C. Fraley and A. E. Raftery. Bayesian regularization for normal mixture estimation and model-based clustering. *Journal of Classification*, 24(2):155–181, 2007.

S. J. Gaffney. *Probabilistic Curve-Aligned Clustering and Prediction with Regression Mixture Models*. PhD thesis, Department of Computer Science, University of California, Irvine, 2004.

Scott Gaffney and Padhraic Smyth. Trajectory clustering with mixtures of regression models. In *Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 63–72. ACM Press, 1999.

T. Hastie and R. Tibshirani. Discriminant analysis by gaussian mixtures. *Journal of the Royal Statistical Society, B*, 58:155–176, 1996.

T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning, Second Edition: Data Mining, Inference, and Prediction*. Springer Series in Statistics. Springer, second edition edition, January 2010.

Trevor Hastie, Andreas Buja, and Robert Tibshirani. Penalized discriminant analysis. *Annals of Statistics*, 23:73–102, 1995.

G. Hébrail, B. Hugueney, Y. Lechevallier, and F. Rossi. Exploratory analysis of functional data via clustering and optimal segmentation. *Neurocomputing*, 73(7-9):1125–1141, March 2010.

DR Hunter and DS Young. Semiparametric mixtures of regressions. *Journal of Nonparametric Statistics*, 24(1):19–38, 2012.

R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3(1):79–87, 1991.

P. N. Jones and G. J. McLachlan. Fitting finite mixture models in a regression context. *Australian Journal of Statistics*, 34(2):233–240, June 1992.

M. I. Jordan and R. A. Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6:181–214, 1994.

Charles Kooperberg and Charles J. Stone. A study of logspline density estimation. *Computational Statistics & Data Analysis*, 12(3):327–347, 1991.

G. J. McLachlan and K. E. Basford. *Mixture Models: Inference and Applications to Clustering*. Marcel Dekker, New York, 1988.

G. J. McLachlan and T. Krishnan. *The EM algorithm and extensions*. New York: Wiley, 1997.

G. J. McLachlan and D. Peel. *Finite mixture models*. New York: Wiley, 2000.

R. E. Quandt. A new approach to estimating switching regressions. *Journal of the American Statistical Association*, 67(338):306–310, 1972.

R. E. Quandt and J. B. Ramsey. Esimating mixtures of normal distributions and switching regressions. *Journal of the American Statistical Association*, 73(364):730–738, 1978.

J. O. Ramsay and B. W. Silverman. *Functional Data Analysis*. Springer Series in Statistics. Springer, June 2005.

Chandan K. Reddy, Hsiao-Dong Chiang, and Bala Rajaratnam. Trust-tech-based expectation maximization for learning finite mixture models. *IEEE Transactions Pattern Analysis and Machine Intelligence*, 30(7):1146–1157, 2008. ISSN 0162-8828.

Sylvia Richardson and Peter J. Green. On Bayesian Analysis of Mixtures with an Unknown Number of Components (with discussion). *Journal of the Royal Statistical Society: Series B (Methodological)*, 59(4):731–792, 1997. ISSN 1467-9868.

M.P. Ruppert, D. Wand and R.J. Carroll. *Semiparametric Regression*. Cambridge University Press, 2003.

G. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6:461–464, 1978.

D. Titterington, A. Smith, and U. Makov. *Statistical Analysis of Finite Mixture Distributions*. John Wiley & Sons, 1985.

Richard D. De Veaux. Mixtures of linear regressions. *Computational Statistics and Data Analysis*, 8(3):227–245, 1989.

K. Viele and B. Tong. Modeling with mixtures of linear regressions. *Statistics and Computing*, 12:315–330, 2002.

Miin-Shen Yang, Chien-Yo Lai, and Chih-Ying Lin. A robust em clustering algorithm for gaussian mixture models. *Pattern Recognition*, 45(11):3950–3961, November 2012. ISSN 0031-3203.

Ka Yee Yeung, Chris Fraley, A. Murua, Adrian E. Raftery, and Walter L. Ruzzo. Model-based clustering and data transformations for gene expression data. *Bioinformatics*, 17(10):977–987, 2001.

DS Young and DR Hunter. Mixtures of regressions with predictor-dependent mixing proportions. *Computational Statistics and Data Analysis*, 55(10):2253–2266, 2010.